Initial State Stabilities and Inverse Engineering in Conflict Resolution

by

Amanda Garcia

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2018

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner	Dr. Takehiro Inohara Professor Department of Social and Human Sciences, School of Environment and Society Tokyo Institute of Technology
Supervisors	Dr. Keith W. Hipel University Professor Department of Systems Design Engineering University of Waterloo
	Dr. Amer Obeidi Adjunct Professor Department of Systems Design Engineering University of Waterloo
Internal Member	Dr. Shi Cao Assistant Professor Department of Systems Design Engineering University of Waterloo
Internal-external Member	Dr. Jatin Nathwani Professor Department of Civil and Environmental Engineering University of Waterloo
Other Member	Dr. Liping Fang Adjunct Professor Department of Systems Design Engineering University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see statement of contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners. I understand that my thesis may be made electronically available to the public.

Statement of Contributions

Chapter 3 is based on the published article (Garcia, Obeidi and Hipel n.d.) which I co-authored with my thesis supervisors Dr. Keith Hipel and Dr. Amer Obeidi.

Chapter 4 is based on two published articles, (Garcia and Hipel 2017) and (Garcia, Obeidi and Hipel 2018). The first article was co-authored with Dr. Keith Hipel and the second was co-authored with Dr. Keith Hipel and Dr. Amer Obeidi.

Chapter 5 is based on a manuscript which has been submitted for publication. It was co-authored with Dr. Amer Obeidi and Dr. Keith Hipel.

Abstract

Two original contributions are made which extend the Graph Model for Conflict Resolution: one is a new family of solution concepts, while the other is a novel methodological approach. In addition to theoretical contributions, applications to complex energy problems are demonstrated; in particular, the consideration of the ongoing Trans Mountain Expansion Project is the first of its kind.

The family of solution concepts, called initial state stabilities, is designed to complement existing solution concepts within the Graph Model framework by modelling both risk-averse and risk-seeking decision-makers. The comparison which underpins these concepts examines the consequences of moving from a given starting state to those of remaining in that state. The types of individuals modelled by these stability concepts represent a new class of decision-makers which, up until now, had not been considered in the Graph Model paradigm.

The innovative methodology presented is designed to "inverse engineer" decision-makers' preferences based on their observable behaviour. The algorithms underlying the inverse engineering methodology are based on the most commonly used stability concepts in the Graph Model for Conflict Resolution and function by reducing the set of possible preference rankings for each decision-maker. The reduction is based on observable moves and counter-moves made by decision-makers. This procedure assists stakeholders in optimizing their own decision-making process based on information gathered about their opponents and can also be used to improve the modelling of strategic interactions.

In addition to providing decision-makers and analysts with up-to-date preference information about opponents, the methodology is also equipped with an ADVICE function which enriches the decision-making process by providing important information regarding potential moves. Decision-makers who use the methods introduced in this thesis are provided with the expected value of each of their possible moves, with the probability of the opponent's next response, and with the opponent reachable states. This insightful data helps establish an accurate picture of the conflict situation and in so doing, aids stakeholders in making strategic decisions. The applicability of this methodology is demonstrated through the study of the conflict surrounding the Trans Mountain Expansion Project in British Columbia, Canada.

Acknowledgements

It goes without saying that this work would not have been possible without the help and encouragement of many individuals. I am extremely fortunate and privileged to have such a supportive network of colleagues, friends, and family.

I am profoundly grateful to my thesis supervisors, Dr. Keith W. Hipel and Dr. Amer Obeidi, for their steadfast support and encouragement throughout my academic journey. From Dr. Hipel I learned of the joy that can accompany research findings and that kindness should always have a place in academia. Dr. Obeidi taught me the importance of finding my own academic voice and of setting ambitious goals. Thank you both for seeing and nurturing my potential.

I would also like to thank my thesis committee: Dr. Takehiro Inohara from the Tokyo Institute of Technology; Dr. Jatin Nathwani from the Department of Civil and Environmental Engineering; Dr. Liping Fang from the Department of Mechanical and Industrial Engineering at Ryerson University; and Dr. Shi Cao from the Department of Systems Design Engineering. Thank you for your insightful feedback which helped improve the quality of this thesis.

I would also like to acknowledge the Natural Sciences and Engineering Research Council of Canada for their financial support of my studies.

Thank you to Dr. Giuseppe Sellaroli for his endless help and patience in developing the programming code for the inverse engineering algorithms. Your coding skills helped make my paper algorithms into a functional reality.

I am grateful to Dr. Stan Dimitrov from the Department of Management Sciences for his guidance during the early stages of the inverse engineering research. I am also thankful to my colleagues in the Department of Systems Design engineering. Janine Blair, Lauren Gratchene, Vicky Lawrence, and Angie Muir offered me excellent support in navigating the bureaucratic process. Thank you to Dr. Matt Borland for encouraging my passion for teaching and to Dr. Deitra Sawh for diving into our teaching journey together. I also appreciate Dr. Carolyn MacGregor for offering her wise counsel.

I am grateful to my colleagues at the Centre for Teaching Excellence who helped me discover and develop my love of teaching: Dr. Donna Ellis, Dr. Svitlana Taraban-Gordon, Dr. Stephanie White, Dr. Kristin Brown, Dr. Crystal Tse, Dr. Lisa Kabesh, and Dr. Tommy Mayberry. Thank you all for supporting my growth as a university teacher.

I would also like to thank each and every one of the students I have had the privilege to teach in some capacity throughout my studies. Thank you for your curiosity, for your kindness, and for challenging me to think (and re-think) deeply about teaching and what it means to me.

I am thankful to my friend and fellow researcher Simone Philpot who started this journey with me and who helped me navigate these uncharted waters. Thank you for your kindred spirit, for your academic street smarts, and for your always sympathetic ear.

Last and certainly not least, I am deeply grateful to my family. To my dad, Alfonso Garcia, who modelled and taught me the virtue of problem-solving and resourcefulness when faced with a challenge; to my mom, Monique Garcia, from whom I learned the values of thinking ahead and asking the right questions; to my sister, Vanessa Garcia, who reminds me not to take myself too seriously and who awes me with her combination of empathy and searing wit rolled into one; to my partner, Giuseppe Sellaroli, who has inexhaustible supplies of patience, humour, and creativity, and who never ceases to amaze me with his support for my hare-brained schemes. Thank you for your unconditional love and support; I love you all.

To women in STEM – past, present, and future

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Strategic interactions among decision-makers (DMs) are common to a wide variety of contexts including economics, evolutionary biology, operations research, negotiations, and military science. Game theory is the study of such interactions; it provides a rich array of mathematical tools to model, analyse, and predict possible resolutions to conflicts. Since its formal introduction in the book by von Neumann and Morgenstern 1944, classical game theory has seen many of its original assumptions re-examined and re-evaluated. The notion of perfect rationality has ceded ground to that of bounded rationality (Simon 1957; Conlisk 1996); a “static" approach has given way to a dynamic, time dependent one (Hofbauer and Sigmund 2003).

Modern approaches are based on evolutionary game theory which combines bounded rationality with dynamics to analyse conflicts and their evolution (Perc et al. 2013; Szabo and Fath 2007). Recent developments in evolutionary game theory include multi-games, in which each DM in a population has different payoffs (Wardil and da Silva 2013; Hashimoto 2014), as well as graph and network approaches in which social networks and spatial games are often represented using graph or lattice structures with edge weights representing utilities (Nowak, Bonhoeffer and May 1994; Roca, Cuesta and Sanches 2009; Galeotti et al. 2010; Perc 2014).

As is detailed in Chapter 2, the Graph Model for Conflict Resolution (Graph Model, for short) is a game-theoretic methodology which provides insights based on a conflict's main DMs, their respective options which are used to define the set of conflict states, and their respective preferences over the set of states. The Graph Model provides the bedrock for this research, which interrogates and examines some of its fundamental assumptions.

1.1 Motivation

This research is motivated by two main questions. The first, “What are other ways to define stability for a state?", looks beyond the typical solution concepts used in the Graph Model. While they provide valuable information and analysis, the most commonly used solution concepts [1] are calculated based on the same comparison of the state of interest to the state(s) resulting from a sequence of DM moves and countermoves. Given this observation, the introduction of diversity in solution concept calculation is particularly relevant as distinct solution concepts

[1] The commonly used solution concepts are Nash, General Metarational, Symmetric Metarational, and Sequential stabilities, which are explained in more detail in Chapter 2.

personify different types of DMs. A thorough Graph Model analysis of a conflict would do well to model a variety of DM types in order to provide valuable insights. Enriching the suite of solution concepts available to users of the Graph Model thus requires conceptualizing a novel state comparison method.

The second question asks “What can one do if one has no preference information about DMs?". As users of the Graph Model know, preferences are a key ingredient in building a credible model of the conflict; however, preference information for DMs may be difficult to come by. Opponents may not want to divulge their preferences to other DMs or to analysts in order to preserve a strategic advantage. For those studying historial conflicts, finding sufficient information to build an accurate preference picture of those involved may be difficult to achieve.

1.2 Research Objectives

In light of the two motivating questions highlighted above, the research objectives can be divided into two groups, each corresponding to one of the questions. The first group of objectives is motivated by the goal of expanding the list of existing solution concepts:

- Determine an alternate method of comparing states
- Based on the above comparison method, develop a family of solution concepts to model a new type of DM

The next group of objectives is designed to address the second question regarding the lack of preference information. The approach “inverse engineers" a DM’s preferences based on their observable moves, and has the following objectives:

- Develop a methodology and algorithms to determine a DM’s preferences based on the behaviour they exhibit
- Enhance the above methodology to provide information which is useful for real-time decision-making

An objective shared by both questions is to demonstrate the validity and applicability of the new solution concepts and methods using real-world case studies.

1.3 Organization of Thesis

This thesis is organized into six chapters, as shown in Figure 1.1. The introductory chapter, Chapter 1, provides the motivation for the research and outlines the research objectives. Chapter 2 gives an overview of the historical development of the Graph Model and details the two main stages which constitute the application of the Graph Model.

The first motivating question (denoted by Question 1 in Figure 1.1) is addressed in Chapter 3, which introduces a new family of solution concepts, called initial state stabilities, whose calculations differ from typical Graph Model solution concepts. Theoretical results linking the new solution concepts to Graph Model concepts are stated and proved. Finally, a case study is used to demonstrate how these solution concepts are applied to real-world disputes and how they compliment existing solution concepts.

The second motivating question (denoted by Question 2 in Figure 1.1) is explored in Chapters 4 and 5. Chapter 4 proposes to answer the question via an innovative methodology and algorithms designed to inverse engineer a DM's preferences based on their observable actions. Beginning with a philosophical reflection on the nature of the approach, the chapter then moves on to review some of the relevant literature regarding the study of preferences in the Graph Model. Next, the necessary concepts underpinning the algorithms are explained, following which the algorithms and the ADVICE function are detailed. The final section in the chapter illustrates the previously introduced methodology using a simple 2-DM case study.

Chapter 5 applies the methods from Chapter 4 to an ongoing pipeline construction conflict which is currently taking place in Alberta, Canada. The conflict, its main DMs, and points of contention are described, following which the inverse engineering model is constructed. Based on the results of the analysis, the strategic insights imparted by the new methodology are emphasized. Finally, Chapter 6 summarizes the main contributions of this work and suggests new research avenues.

Figure 1.1: Organization of the thesis

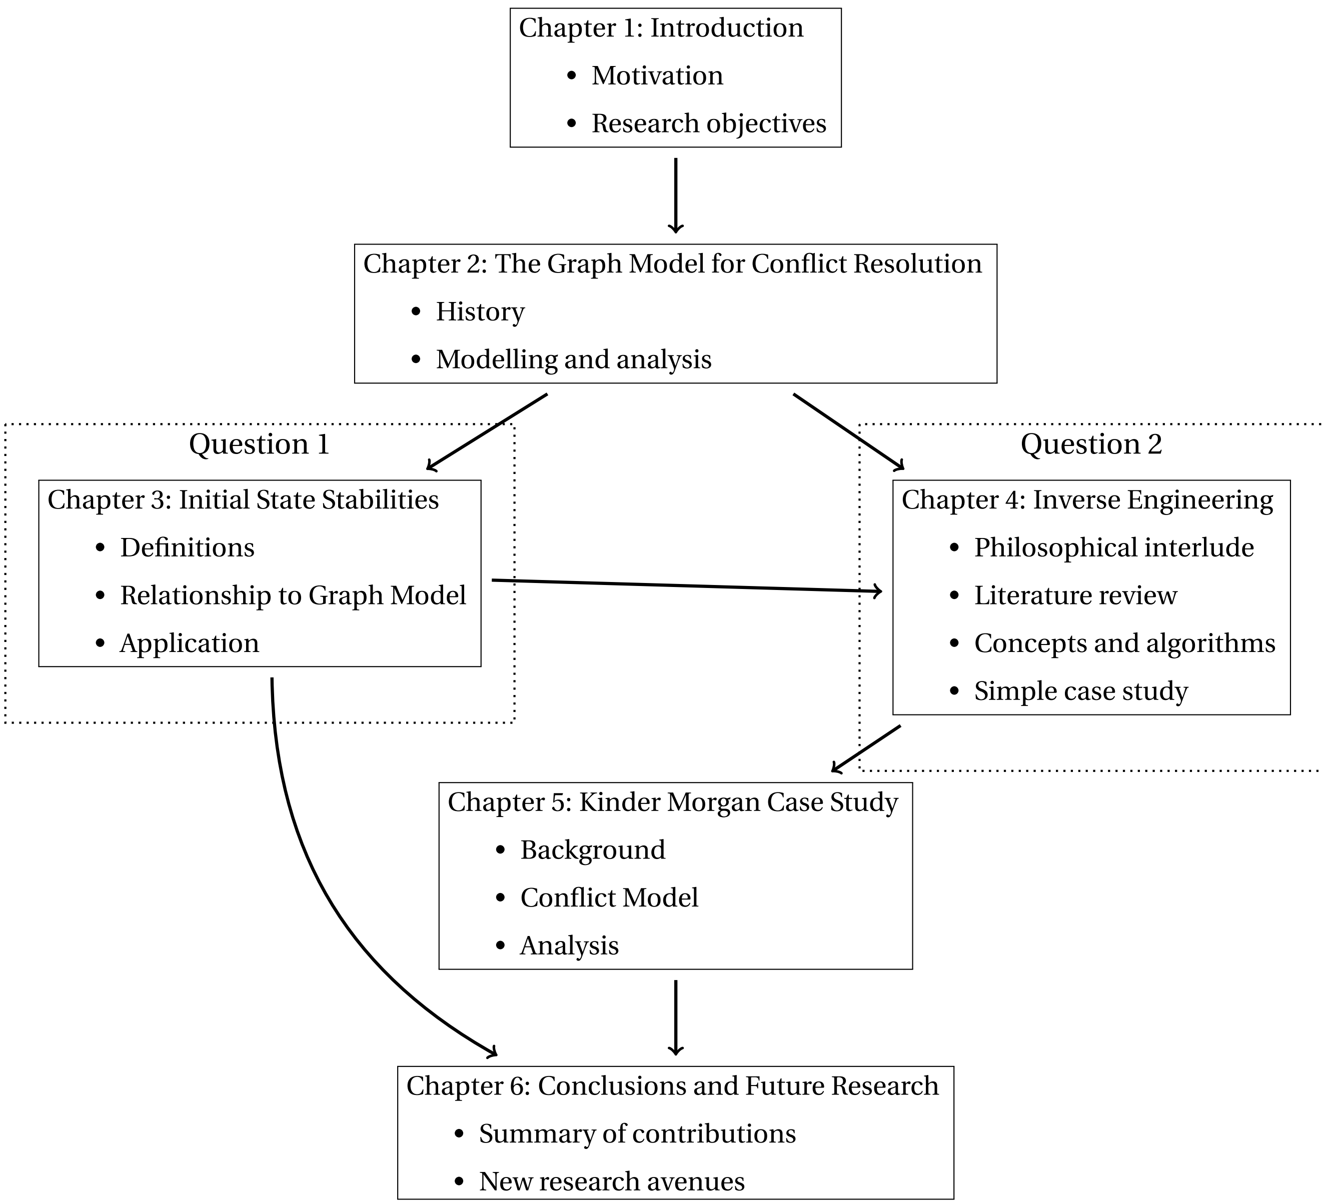

Chapter 2

The Graph Model for Conflict Resolution

Traditionally, DM preferences are modelled with utility functions. Such functions are meant to capture the benefit (utility) that each DM derives from a particular state or outcome; many approaches seek to maximise the utility functions of the DM(s) in question. Cardinal utilities and preferences can, however, be replaced by ordinal preferences in which states are pairwise ranked and compared.

Relative preferences have some advantages over cardinal ones. First, relative preference statements such as "I prefer tea to coffee" are more reflective of how people think of and describe their own preferences; one is unlikely to encounter an individual who would state "I prefer tea 2.8 times to coffee". By noting only the DM's ranking of items (tea > coffee) rather than the ranking and a metric (tea = 2.8 coffee), ordinal preferences avoid the problems of trying to determine a precise utility for each item and of performing interpersonal comparisons of utility. Such comparisons are by nature quite difficult since it is not clear what the "exchange rate" is between one individual's unit of utility and another's.

Second, relative preferences allow for intransitivity in preference statements, a common occurrence which nevertheless is precluded by utility functions since modelling preferences with utility functions requires that they be transitive, complete, and monotone. Intransitive preferences occur when three or more items or statements are being compared and "cycles" in preferences occur. For example, an individual may prefer tea to coffee, coffee to juice, but juice to tea.

The Graph Model is a well-known methodology for analysing complex conflicts involving several DMs which uses ordinal preferences. This methodology analyses DM moves and counter-moves in order to provide strategic insights and to predict possible resolutions of the conflict under study.

2.1 A Brief History of the Graph Model

The Graph Model finds its roots in metagame theory which was pioneered by Nigel Howard (Howard 1971). This non-quantitative approach is based on the analysis of options at the disposal of stakeholders and is meant as a more practical and intuitive answer to the rigidity of classical game theory. Rather than working with utility functions and cardinal preferences, metagame analysis focuses on the possible outcomes of a conflict and on which DMs can influence them.

Further work in this field by Fraser, Hipel, Kilgour, and Fang resulted in the development of conflict analysis (Fraser and Hipel 1979; Fraser and Hipel 1984) and, eventually, of the Graph Model (Kilgour, Hipel and Fang 1987; Fang, Hipel and Kilgour 1993). These techniques expanded the scope and applicability of metagame theory by developing a systematic, operationalised way to analyse conflicts and by introducing several new and valuable solution concepts. A decision support system (DSS), GMCR II, was eventually developed to facilitate Graph Model calculations (Fang et al. 2003b; Fang et al. 2003a).

Parallel progress in metagame analysis also included the development of drama theory, which uses the metaphor of theatre to analyse the dynamics of multi-DM conflicts. The conflicts, viewed through the lens of theatre, undergo phases of scene-setting, build-up, climax, decision, and denouement (Howard 1999; Bryant 2015).

Since the introduction of the Graph Model in the late 1980s, many enhancements and improvements have been made. Notable examples include the matrix formulation of the Graph Model which uses matrix operations to greatly simplify stability calculations (Xu, Hipel and Kilgour 2009a; Xu, Li et al. 2009) and coalition analysis which examines potential resolutions based on cooperation among DMs (Inohara and Hipel 2008)[1].

The Graph Model's robust and relevant results, coupled with its ease of use with decision support systems such as GMCR II and GMCR + (Kinsara, Petersons et al. 2015), have made it a popular methodology to study a wide range of conflicts. Most recently, the Graph Model has been used to analyse energy conflicts (Garcia, Obeidi and Hipel 2016; Matbouli, Hipel and Kilgour 2014); water conflicts (Hipel, Kilgour and Kinsara 2014; Philpot, Hipel and Johnson 2016; He, Kilgour and Hipel 2016); and environmental disputes (Bashar, Hipel and Kilgour 2012; Madani 2013). Advancing and improving the Graph Model remains a lively, multi-disciplinary research area to this day.

2.2 Procedure

The Graph Model procedure is executed in two broad steps: modelling and analysis. The general Graph Model procedure is illustrated in Figure 2.1, which is adapted from (Fang, Hipel and Kilgour 1993; Xu, Hipel, Kilgour and Fang 2018). In the next sections, the procedures are reviewed and essential definitions and notation are provided.

2.2.1 Modelling

In the modelling phase, the relevant conflict parameters are determined and/or calculated. First, the DMs are identified. DMs are individuals or groups with decision-making power and who have a direct influence on the conflict. The set of DMs is denoted by $N = \{1, 2, \ldots, n\}$ with n being the total number of DMs and each DM being arbitrarily assigned a unique numerical identifier.

Next, the set of options for each DM is identified. Options are binary; by convention, an option which has been chosen is denoted by a 1, while an option which has not been selected

[1] Graph Model extensions pertinent to the work presented in this thesis will be detailed in Chapter 4.

Figure 2.1: The Graph Model Procedure

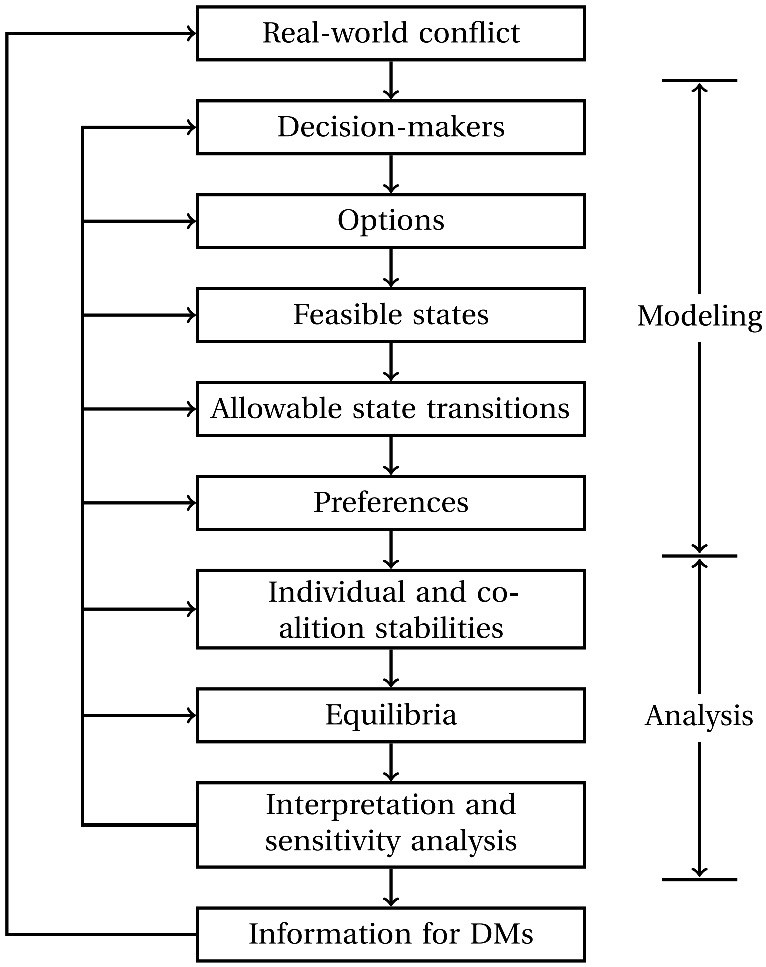

is denoted by a 0. Each DM $i \in N$ has a set O_i of options, with o_{ij} denoting DM i's jth option. The set of all options in a conflict is $O = \cup_{i \in N} O_i$. During the conflict, DMs move among their strategies, selecting and un-selecting the options at their disposal.

Definition 2.1. A *strategy* for DM i is a mapping $g : O_i \rightarrow \{0, 1\}$ such that
$$g(o_{ij}) = \begin{cases} 1 & \text{if DM } i \text{ selects option } o_{ij} \\ 0 & \text{otherwise} \end{cases}$$

A DM's strategy is simply an assignment of 0/1 to each of that DM's options. Combining fixed strategies from all of the DMs results in a conflict state.

Definition 2.2. A *state* is a mapping $f : O \rightarrow \{0, 1\}$ such that
$$f(o_{ij}) = \begin{cases} 1 & \text{if DM } i \text{ selects option } o_{ij} \text{ for } i = 1, 2, \ldots n \\ 0 & \text{otherwise} \end{cases}$$

Table 2.1: Types of Infeasible States

Type of infeasibility	**Example**
Logical infeasibility for one DM	Mutually exclusive options: building a project and not building a project
Preferential infeasibility for one DM	States which involve a strategy selection by a DM which they would not be expected to take: a review commission choosing not to make recommendations
Logical infeasibility among a subset of DMs	DMs vying for an indivisible resource: a house in a divorce settlement
Preferential infeasibility among a subset of DMs	Preferentially infeasible outcomes for at least one DM which involve strategy choices by at least two DMs: a DM will let an opponent commit a legal infraction without taking legal action

A conflict with $|O| = m$ options has 2^m possible states; however, not all of these may be feasible. Removing infeasible states is thus the next step in the modelling phase. A state may be removed for one of four reasons; the types of infeasible states are summarized in Table 2.1

Once infeasible states have been removed, only feasible states remain. The set of feasible states is denoted by $S = \{1, 2, \ldots, s\}$ and represents outcomes of the conflict under study that could realistically occur. It is quite common for conflicts having a large number of possible states to be reduced to less than 100 (often less than 50) feasible states.

Next, allowable state transitions are determined. Moves among states may be either reversible or irreversible. In the former case, a move to state $s_k \in S$ can be done or undone; in the latter case, a move to state $s_j \in S$ cannot be undone.

The final modelling step consists in developing a preference ranking for each DM; that is, a ranking of the feasible states from most to least preferred. Preference rankings can be determined through direct consultation with DMs or through historical investigation as necessary.

A simple, binary preference structure is used to express preference rankings in the standard Graph Model[2]: the relation $s_k \succ_i s_j$ indicates that state s_k is preferred by DM i to state s_j, while $s_k \sim_i s_j$ means that states s_k and s_j are equally preferred by DM i. The set of relations $\{\prec_i, \sim_i\}$ is complete: any two states can be compared by a DM who will either prefer one over the other or prefer them equally.[3]

One of the advantages of the Graph Model lies in its capacity to accommodate a rich set of preference types, including both transitive and intransitive preferences. In the former case, this means that the ranking satisfies the assumption of transitivity which states that if $s_k \prec_i s_j$ and $s_j \prec_i s_l$, then it must be that $s_k \prec_i s_l$. In the latter case, a DM's preference ranking does

[2] Many other preference structures have been developed and will be discussed in Chapter 4.

[3] In this context, "or" is understood as an exclusive or: the DM cannot affirm both or neither of these statements.

not satisfy the above condition. It is therefore possible to have preference “cycles" such that $s_k \prec_i s_j$, $s_j \prec_i s_l$, and $s_l \prec_i s_k$.

Given the DMs, their options, the set of feasible states, allowable transitions, and DMs' preferences, the Graph Model of the conflict can be formed. The conflict is modelled as a set of finite directed graphs $D_i = (S, A_i)$ with $i \in N$ where the set of vertices is given by the states and an arc $a \in A_i$ between vertices s_k and s_j indicates that DM i can unilaterally move from state s_k to state s_j in one step. For an n-DM conflict with $n \geq 2$, the Graph Model illustrates the conflict as a set of finite directed graphs in which vertices represent feasible states and arcs depict moves controlled by DMs.

2.2.2 Analysis

Once the model is formed, one can proceed to the analysis phase. Here, the information provided in the modelling phase is used to determine likely outcomes of the conflict as well as its possible evolution from a status quo state. Broadly, states are examined for stability at the individual DM level, at the coalition level, and at the conflict level using a variety of solution concepts. A stable state is one from which there is no incentive to move away; states which are stable under a given solution concept for all DMs are called equilibrium states and represent possible resolutions of the conflict.[4]

Solution concepts (also commonly referred to as stability concepts) represent different behavioural profiles that could apply to the DMs. The choice to use one solution concept over another depends on whether the focal DM is willing to accept risk, has knowledge of their opponents' preferences, and how much foresight is to be considered. Analysts and DMs may use one or more solution concepts when analysing a conflict to examine how the choice of solution concept can affect the equilibrium results.

In order to define the solution concepts used in the Graph Model, some preliminary notations and definitions are required. First, the sets of unilateral moves and unilateral improvements:

Definition 2.3. The set of *unilateral moves* (UMs) for DM i from state $s_k \in S$ is given by $R_i(s_k) = \{s_j : s_k s_j \in A_i\}$. In words, $R_i(s_k)$ denotes the set of states that DM i can unilaterally reach in one step from state $s_k \in S$.

Definition 2.4. The set of *unilateral improvements* (UIs) for DM i from state $s_k \in S$ is given by $R_i^+(s_k) = \{s_j : s_j \in R_i(s_k) \text{ and } s_k \prec_i s_j\}$. In words, $R_i^+(s_k)$ denotes the set of states that DM i can unilaterally reach in one step from state $s_k \in S$ and that are more preferred to state s_k.

UMs and UIs provide information on which states are reachable from a given starting state by each DM and, in the case of UIs, which of these reachable states are improvements from the starting state. UMs and UIs are the basic DM movements which are at the heart of solution concepts.

In n-DM conflicts, one can define UMs for subsets of DMs. A non-empty subset $H \subseteq N$ is called a *coalition*. Coalition UMs are defined as sequences of moves in which DMs in the coalition may move more than once, but not twice consecutively (Fang, Hipel and Kilgour

[4] The term "resolutions" is used in the broad sense of "outcomes"; a resolution to a conflict need not please all DMs.

1993; Xu, Kilgour et al. 2014). The notation $\Omega_H(s, s_1) \subseteq N$ denotes the set of all last DMs in legal sequences from s to s_1.

Definition 2.5. For $s \in S$, $H \subseteq N$ a non-empty subset of DMs, and assuming $\Omega_H(s, s_1) = \emptyset$ to start, a *unilateral move by coalition* H is a member of $R_H(s) \subseteq S$, defined inductively by:

(1) if $j \in H$ and $s_1 \in R_j(s)$, then $s_1 \in R_H(s)$, and $\Omega_H(s, s_1) = \Omega_H(s, s_1) \cup \{j\}$;

(2) if $j \in H$, $s_1 \in R_H(s)$, $s_2 \in R_j(s_1)$, then provided $\Omega_H(s, s_1) \neq \{j\}$, $s_2 \in R_H(s)$ and $\Omega_H(s, s_2) = \Omega_H(s, s_2) \cup \{j\}$.

According to (1) in Definition 2.5, all the reachable states from s are identified and added to $R_H(s)$ and DM j is added to $\Omega_H(s, s_1)$, the set of all last DMs in legal sequences from s to s_1. At (2), all states reachable from the states identified in (1) are identified and also added to $R_H(s)$ provided that the DMs moving to these states are not the only DMs in the set of last movers $\Omega_H(s, s_1)$. The process is repeated from (2) until there are no more states to add to $R_H(s)$. The set $R_H(s)$ thus contains the states that the conflict could reach from the starting state s if only the DMs in H are allowed to move. The definition for UIs for sets of DMs is analogous; $R(s)$ is simply replaced by $R^+(s)$:

Definition 2.6. For $s \in S$, $H \subseteq N$ a non-empty subset of DMs, and assuming $\Omega^+_H(s, s_1) = \emptyset$ to start, a *unilateral improvement by coalition* H is a member of $R^+_H(s) \subseteq S$, defined inductively by:

(1) if $j \in H$ and $s_1 \in R^+_j(s)$, then $s_1 \in R^+_H(s)$, and $\Omega^+_H(s, s_1) = \Omega^+_H(s, s_1) \cup \{j\}$;

(2) if $j \in H$, $s_1 \in R_H(s)$, $s_2 \in R_j(s_1)$, then provided $\Omega^+_H(s, s_1) \neq \{j\}$, $s_2 \in R_H(s)$ and $\Omega^+_H(s, s_2) = \Omega^+_H(s, s_2) \cup \{j\}$.

The notions of states, preferences, UMs, and UIs are used to define several solution concepts in the Graph Model, the most common of which are summarized in Table 2.2 and defined below. In the following definitions, the notation $N - i$ denotes the set of DMs excluding DM i.

Definition 2.7. A state $s_k \in S$ is *Nash stable* for DM $i \in N$ if and only if (iff) $R^+_i(s_k) = \emptyset$. In words, a state is Nash stable (R) if there is no available UI for the DM (Fraser and Hipel 1979).

At a Nash stable state, the DM has no incentive to move away since there are no available UIs. The Nash stable state is the best state that the DM can achieve, assuming that the opponents' strategies remain fixed.

Definition 2.8. A state $s_k \in S$ is *general metarational (GMR) stable* for DM $i \in N$ iff for every $s \in R^+_i(s_k)$ there exists at least one state $s_x \in R_{N-i}(s)$ with $s_x \precsim_i s_k$ (Fraser and Hipel 1979) .

In words, a state s_k is GMR stable for DM i if for every UI available to DM i, there is a state in the coalition UMs by the other DMs which is less than or equally preferred by i to s_k. DM i is guarding against the possibility that moving to a UI could potentially result in a worse outcome than simply staying put.

Definition 2.9. A state $s_k \in S$ is *sequentially (SEQ) stable* for DM $i \in N$ iff for every $s \in R^+_i(s_k)$ there exists at least one state $s_x \in R^+_{N-i}(s)$ with $s_x \precsim_i s_k$ (Fraser and Hipel 1979; Fraser and Hipel 1984).

A state s_k is SEQ stable for DM i if for every UI available to DM i, there is a state in the coalition UIs by the other DMs which is less than or equally preferred by i to s_k. SEQ stability and GMR stability are similar but for the fact that SEQ stability requires knowledge of the opponents' preferences. While GMR stability is concerned with opponent UMs, SEQ stability deals with opponent UIs.

Definition 2.10. A state $s_k \in S$ is *symmetric metarational (SMR) stable* for DM $i \in N$ iff for every $s \in R_i^+(s_k)$ there exists at least one state $s_x \in R_{N-i}(s)$ with $s_x \precsim_i s_k$ and for all $s_y \in R_i(s_x)$, $s_y \precsim_i s_k$ (Fraser and Hipel 1979).

In other words, a state s_k is SMR stable for DM i if for every UI available to DM i, there is a state s_x in the coalition UM by the other DMs which is less than or equally preferred by i to s_k and any UM by i from s_x is less than or equally preferred by i to s_k. DM i is anticipating that it will have the opportunity to respond to the opponents' counter-move. At SMR stable states, DM i's response still does not result in escaping the sanction which could potentially be imposed by the opponents.

Table 2.2: Summary of Graph Model solution concepts for a state $s_k \in S$ for DM $i \in N$

Solution Concept	**Definition**
Nash Stability (R)	$R_i^+(s_k) = \emptyset$.
General Metarational (GMR)	$\forall s \in R_i^+(s_k)\ \exists s_x \in R_{N-i}(s)$ with $s_x \precsim_i s_k$.
Sequential Stability (SEQ)	$\forall s \in R_i^+(s_k)\ \exists s_x \in R_{N-i}^+(s)$ with $s_x \precsim_i s_k$.
Symmetric Metarational (SMR)	$\forall s \in R_i^+(s_k)\ \exists s_x \in R_{N-i}(s)$ with $s_x \precsim_i s_k$ and $\forall s_y \in R_i(s_x)$, $s_y \precsim_i s_k$.

Other solution concepts such as limited move stability (Fang, Hipel and Kilgour 1993), non-myopic stability (Brams and Wittman 1981), and symmetric sequential stability (Rego and Vieira 2017) exist; however, the most commonly used solution concepts in a Graph Model analysis are Nash, GMR, SMR, and SEQ stabilities. Table 2.3, adapted from (Kilgour, Hipel and Fang 1987), describes important characteristics of each of the solution concepts presented above and also shows the different behavioural profiles associated to each solution concept.

One can see that Nash stability looks only one step ahead and is risk averse in the sense that it does not consider possible opponent counter-moves but focuses on moving to a state which cannot be improved upon at the point in time of interest. Next, GMR and SEQ stabilities look two steps into the future: the DM's initial move to a UI and possible opponent counter-moves. For GMR stability, no knowledge of opponent preferences is assumed; equivalently, one could assume that opponents may be willing to accept a disimprovement in order to punish the DM. SEQ stability, however, requires knowledge of the opponent's preferences and assumes that opponents will not accept disimprovements. Finally, SMR stability looks yet another step ahead by allowing the DM to respond to opponent counter-moves.

Once stabilities have been calculated for each individual DM, conflict equilibria can be found. Recall that a state which is stable under a particular solution concept for all DMs is an *equilibrium* for that solution concept. Thus, a Nash equilibrium is a state which is Nash stable for all DMs and similarly for the other solution concepts. Equilibrium states provide DMs and analysts with possible resolutions of the conflict.

Stabilities are also calculated for coalitions of DMs. *Coalition analysis* provides additional insights on possible resolutions by examining whether DMs could cooperate and, in so doing, reach a more preferred resolution than could be reached if they operated individually (Xu, Kilgour et al. 2014). This analysis also reveals which equilibria are susceptible to shifts (sometimes called equilibrium jumps) by coalitions of DMs.

It is also common to perform *sensitivity analysis* in which conflict parameters are slightly perturbed to determine the robustness of the initial analysis results. Sensitivity analysis examines "What if?" scenarios and allows one to see how sensitive the equilibria are to changes in DMs, preferences, and options. Once analysis has been performed, the information and insights are used to help DMs and analysts make better decisions about the conflict situation.

Table 2.3: Solution Concepts and Human Behaviour

Solution Concept	**Description**	**Foresight**	**Knowledge of Preferences**	**Disimprovement**	**Risk**
Nash	DM cannot unilaterally move to a more preferred state	Low	Own	Never	Ignores risk
GMR	All the DM's UIs are sanctioned by opponent UMs	Medium	Own	By opponents	Conservative
SEQ	All the DM's UIs are sanctioned by opponent UIs	Medium	All	Never	Some risk
SMR	All the DM's UIs are sanctioned by opponent UMs, even after the DM responds	Medium	Own	By opponents	Conservative

2.3 Extensions to the Graph Model

Beyond the basic Graph Model procedure discussed in Section 2.2, many new extensions and approaches have been developed over the years. This section briefly summarizes some of the extensions to the Graph Model, grouped by topic.

2.3.1 Evolution of a Conflict

Since the Graph Model typically analyses a conflict at a fixed point in time, it is common to track its evolution using methods such as *status quo analysis* (Li, Kilgour and Hipel 2005). Shifts in conflict equilibria due to changes in the DMs' preferences are studied using *robustness of equilibria* methods (Matbouli, Kilgour and Hipel 2015). The length of time that a DM remains in a given state can be examined using a *state transition time* analysis (Inohara 2016).

2.3.2 Inverse GMCR

As its name implies, this approach reverses some of the main steps in the Graph Model. Rather than predicting equilibrium states, the inverse GMCR begins by selecting the desired resolution, then determines which preferences are necessary to achieve it (Kinsara, Kilgour and Hipel 2014; Kinsara, Petersons et al. 2015; Wang et al. 2018).

2.3.3 Power

Power imbalances and asymmetries often occur in conflicts. *Hierarchical* graph models are designed to study DMs which are involved in more than one dispute and may have differing levels of power in each (He, Kilgour, Hipel and Bashar 2013; He, Kilgour and Hipel 2016). The study of *power asymmetry* in a conflict examines how a more powerful DM may be able to influence the preferences of the others by virtue of its privileged position (Yu et al. 2015).

2.3.4 Preference Structures

In this category are grouped different approaches to defining the preference relations used by DMs when ranking the conflict states. *Fuzzy preferences* (Bashar, Kilgour and Hipel 2014; Bashar, Hipel and Kilgour 2012), *grey preferences* (Kuang et al. 2015; Han, Nguyen and Xu 2013), *probabilistic preferences* (Rego and dos Santos 2015), and *strength of preferences* (Xu, Hipel and Kilgour 2009b; Hamouda, Kilgour and Hipel 2006) each define new sets of preference relations based on their chosen mathematical structure. Rather than the standard set $\{\succ_i, \sim_i\}$, a more nuanced approach which allows for degrees of preference is introduced.

2.3.5 Preference Elicitation

These extensions focus on ways in which to develop preference rankings for DMs. *Multi-criteria* methods (Ke, Fu et al. 2012) use existing multi-criteria techniques and apply them to the Graph Model framework. More recently, *value-focused* approaches propose ways to link a DM's values with their preference ranking (Bristow, Fang and Hipel 2014; Philpot, Hipel and Johnson 2017).

2.3.6 Psychological Factors

Extensions in this vein seek to incorporate psychological factors which can influence decision-making into the Graph Model. Work on *attitudes* examines how a DM's attitude (positive, negative, neutral) toward themselves and other DMs can affect the conflict outcomes (S. B. Walker, Hipel and Inohara 2009; S. Walker, Hipel and Inohara 2012). Accounting for the role of *emotions* in a conflict examines the important role that emotions play in how a DM conceptualizes the dispute (Obeidi, Hipel and Kilgour 2005; Obeidi, Kilgour and Hipel 2005). Finally, *hypergames* model how misperceptions with respect to the DMs, the conflict states, preferences, or any combination of the above can affect a conflict's outcome (Aljefri, Hipel and Fang 2018).

2.3.7 Matrix Methods

The *matrix representation* of the Graph Model uses matrices and their operations to both represent a conflict and to calculate its equilibria (Xu, Li et al. 2009; Xu, Hipel and Kilgour 2009a). Over the years, many of the extensions previously mentioned, including coalition analysis (Xu, Kilgour et al. 2014), attitudes (S. B. Walker, Hipel and Xu 2013), uncertain preferences (Xu, Hipel, Kilgour and Chen 2010), grey preferences (Han, Nguyen and Xu 2013), and the inverse Graph Model (Wang et al. 2018) have been expressed using a matrix representation.

2.4 Chapter Summary

This chapter reviewed the Graph Model's historical origins and outlined the Graph Model procedure, which is executed in two phases. The modelling phase consist in selecting the DMs, their options, removing infeasible states, and formulating preference rankings for each DM. In the analysis phase, both individual and coalition stabilities are determined, the equilibria are calculated, followed by sensitivity analysis and interpretation.

Finally, many of the extensions to the Graph Model were summarised. The extensions were grouped by topic and included the study of the evolution of a conflict, inverse GMCR, the role of power, diversity in preference structures and preference elicitation, psychological factors, and matrix methods.

Chapter 3

Initial State Stabilities

As discussed in the previous chapter, the Graph Model requires three "ingredients" to model a conflict: DMs, each DM's options, and each DM's preferences over the set of states. Using this information, one is able to analyse the conflict and find stable states; that is, states from which a DM does not want to move away. The stability of a particular state is always determined with respect to a particular solution concept (also called stability concept). Solution concepts are designed to mathematically express how DMs, with varying levels of risk aversion and foresight, behave in strategic interactions. A standard application of the Graph Model methodology to a conflict typically verifies four types of solution concepts: Nash, GMR, SEQ, and SMR (see Chapter 2 for details on these solution concepts).

A common research approach within the Graph Model is to apply a new preference structure and derive the corresponding version of the four core solution concepts (see Bashar, Obeidi et al. 2016 for an example). Although these approaches do lead to new solution concepts, they remain variations on the core solution concepts and as such retain their general properties with respect to risk aversion and foresight.

Research into developing new Graph Model solution concepts has generally been limited. A decade ago, a series of papers introduced a generalization of general metarational and sequential stabilities (Zeng et al. 2005; Zeng et al. 2006; Zeng et al. 2007). More recently, a new solution concept, symmetric sequential stability, was introduced to expand the foresight of sequential stability (Rego and Vieira 2017). Both of these efforts generated new solution concepts which grew from the core stability concept definitions.

As the literature suggests, analysing a conflict with an enriched set of solution concepts can yield more valuable results to analysts, who gain additional insights into possible conflict resolutions. Since solution concepts are meant to simulate DM behaviour, states which are equilibria under a large number of solution concepts (i.e., under a large number of possible behaviours) may be more likely resolutions of the conflict (Madani 2013; Madani and Hipel 2011; Kilgour and Eden 2010; Kilgour, Hipel, Fang and Peng 2001). The work presented in this chapter thus seeks to complement and expand the range of solutions concepts currently available within the Graph Model framework.

A family of four new solution concepts based on the concept of initial states is introduced. The types of individuals modelled by these stability concepts represent a new class of DM which, up until now, had not been considered in the Graph Model paradigm. As will be shown,

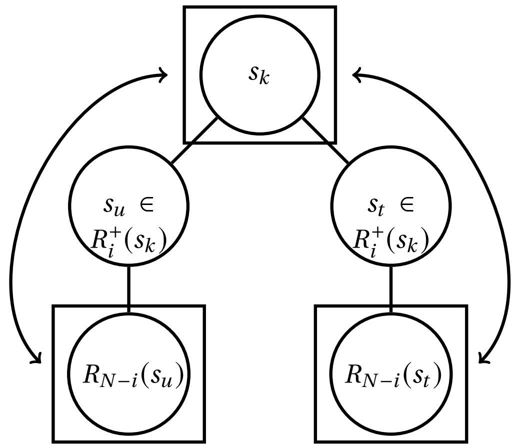

Figure 3.1: Comparison for Graph Model stabilities

the considerations and calculations made with the new solution concepts reflect a different comparison process.

3.1 Stabilities in the Graph Model

The most commonly used Graph Model solution concepts which take into account opponent moves (GMR, SMR, and SEQ) calculate stability in a similar manner: the general idea is to compare some starting state s_k to states in $R_{N-i}(s)$ (or in $R^+_{N-i}(s)$) with $s \in R^+_i(s_k)$. If for all $s \in R^+_i(s_k)$ any $s_x \in R_{N-i}(s)$ (or $s_x \in R^+_{N-i}(s)$) is less preferred to s_k, then the state is said to be stable[1].

3.1.1 Description

Figure 3.1 illustrates how these comparisons work. At the top is the starting state s_k for which stability is being calculated. From s_k, DM i has UIs s_t and s_u[2]. Next, the opponent's moves (either UMs or UIs depending on the stability concept being applied) from DM i's UIs are calculated, giving sets $R_{N-i}(s_t)$ and $R_{N-i}(s_u)$, respectively. The states in $R_{N-i}(s_t)$ and in $R_{N-i}(s_u)$ are then compared to s_k; if at least one state in $R_{N-i}(s_t)$ and at least one state in $R_{N-i}(s_u)$ is less preferred to s_k, then s_k is stable. As shown with the boxes in the figure, the comparison is being made between s_k and $R_{N-i}(s_t)$ and between s_k and $R_{N-i}(s_u)$.

The underlining motive for DMs' moves in this stability analysis is utility maximization based on (at most three) moves by DMs. Thus, the idea that a DM might remain at s_k is only considered when s_k is more preferred to any of the states in $R_{N-i}(s_t)$ and $R_{N-i}(s_u)$. However, there remains the interesting question of the downstream implications of DM i deciding to remain at s_k – could DM i benefit from staying or could there be sanctions compared to taking an available UI? This line of inquiry motivates the introduction of new family of stability

[1] SMR stability allows DM i to move once more before checking for stability.

[2] There may be more or less than two possible UIs; the number of UIs depicted in the figure is for illustration purposes only.

concepts which compares parallel scenarios in which DM i stays at s_k to those in which DM i moves to an available UI.

3.1.2 Example

By way of example, a dispute that occurred between the Ministry of Environment (MoE), the local government (LG), and Uniroyal Chemical Ltd (UR) in the town of Elmira, Ontario is considered. This conflict was first introduced in Hipel, Fang et al. 1993 and has been used to study a variety of Graph Model extensions.

In 1989, the MoE discovered that the town of Elmira's main source of water, an underground aquifer, was contaminated with N-nitroso demethylamine (NDMA). It was immediately suspected that UR, which made pesticides and rubber products, was to blame for the presence of the carcinogen in the aquifer. The MoE issued a Control Order (CO) requiring that UR correct the situation; however, UR rejected the order and launched an appeal. Local governments from the Regional Municipality of Waterloo and the Township of Woolwich banded together in order to safeguard their interests and attempted to pressure UR into accepting the CO. At the point in time of analysis (1991), UR was appealing the CO and the conflict had stagnated.

Table 3.1 illustrates the DMs, their options, and the feasible states for this conflict derived in Hipel, Fang et al. 1993. The MoE has the option to modify the CO issued by making it more favourable to UR. UR can delay the appeal process, accept the CO being offered by the MoE, or abandon its site in Elmira. Finally, the LG can insist that the original CO be applied. The possible feasible states that can be formed in this conflict are shown as columns of 1s and 0s in this table where a "1" means that the option opposite has been selected by the DM controlling it and a "0" indicates that it was not chosen. The dashes in the last column convey that the DM's choice of option is irrelevant; this occurs because UR's decision to abandon the Elmira site renders the rest of the DM's choices moot. For convenience, the states are labelled s_1 to s_9 in the bottom row of the table. The format used in Table 3.1 is called Option Form and was originally introduced by Howard (Howard 1971); it is commonly used within the Graph Model to represent states in a realistic fashion.

Table 3.1: DMs, options, and feasible states for the Elmira conflict

Ministry of Environment (MoE)									
Modify CO favourably	0	1	0	1	0	1	0	1	-
Uniroyal (UR)									
Delay the appeal	1	1	0	0	1	1	0	0	-
Accept CO being offered	0	0	1	1	0	0	1	1	-
Abandon the site	0	0	0	0	0	0	0	0	1
Local Government (LG)									
Insist on original CO	0	0	0	0	1	1	1	1	-
States	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9

Given the feasible states in the conflict, each DM has its own ranking of the states depending on their goals. The MoE would like the aquifer to be cleaned, UR prefers that the CO be either modified or rescinded, and the LG wants to protects its citizens. The Graph Model representation of this conflict is shown in Figure 3.2. The dashed lines represent the moves available to the MoE; the solid lines correspond to moves available to UR, and the dotted lines represent the moves available to the LG. The preference rankings for the DMs from the original analysis, with states listed from most preferred on the left to least preferred on the right, are kept.

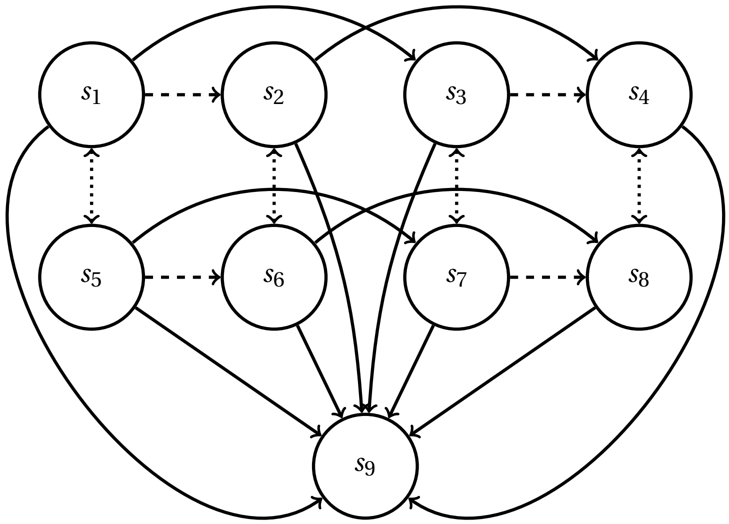

MoE(dashed lines): $s_7 \succ s_3 \succ s_4 \succ s_8 \succ s_5 \succ s_1 \succ s_2 \succ s_6 \succ s_9$
UR (solid lines): $s_1 \succ s_4 \succ s_8 \succ s_5 \succ s_9 \succ s_2 \succ s_3 \succ s_7 \succ s_6$
LG (dotted lines): $s_7 \succ s_3 \succ s_5 \succ s_1 \succ s_8 \succ s_6 \succ s_4 \succ s_2 \succ s_9$

Figure 3.2: Graph Model of the Elmira conflict

3.2 Initial State Stabilities

The general idea of the comparisons performed in the new solution concepts is illustrated in Figure 3.3. From some starting state s_k, DM i can choose to stay at s_k or move to some available UI s_u or s_t[3]. Opponent moves from each of these states ($R_{N-i}(s_k)$, $R_{N-i}(s_u)$, and $R_{N-i}(s_t)$) are calculated and states in $R_{N-i}(s_k)$ are compared to states in $R_{N-i}(s_u)$ and to states in $R_{N-i}(s_t)$. This comparison is assessing outcomes that might occur if DM i remains at s_k to those which might transpire if DM i were to take a UI from state s_k. Since the new stability concepts consider what occurs when DM i remains at some initial state s_k, they will be referred to as *initial state stabilities.*

The main difference between initial state stabilities and the traditional Graph Model stabilities lies in which sets of counter-moves by opponents are under consideration. Graph

[3]The number of available UIs shown in this example is for illustration purposes only.

Model stabilities consider counter-moves from the focal DM's set of UIs; initial state stabilities consider counter-moves from the focal DM's set of UIs *and* from the initial state s_k. Initial state stabilities thus treat the starting state s_k in the same way as Graph Model stabilities treat UIs, that is, by considering which states the opponents can reach from there.

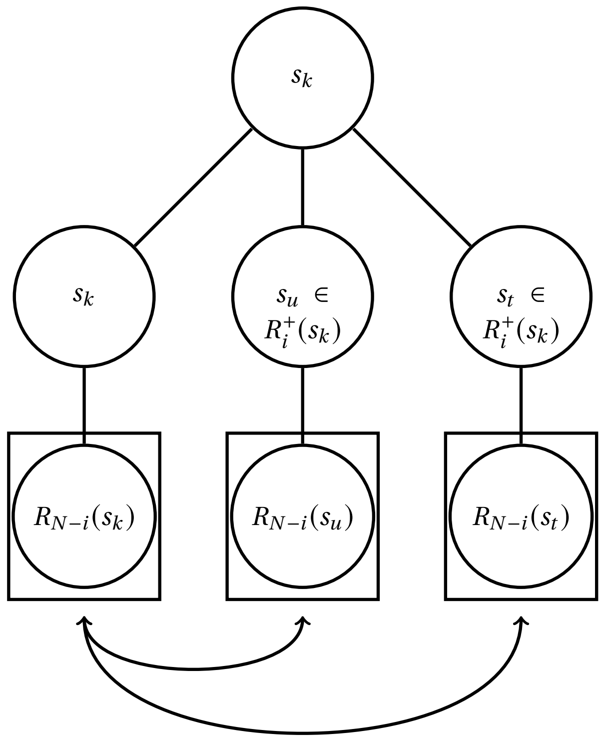

Figure 3.3: Comparison for new stabilities

The question that now arises is how to compare states in $R_{N-i}(s_k)$ to those in $R_{N-i}(s_u)$ and in $R_{N-i}(s_t)$ in a way that provides useful information. To this end, a dichotomous analysis which reflects a DM's level of risk-taking is proposed: “optimistic" and “pessimistic" initial state stabilities. For the present purposes, the terms “optimistic" and “pessimistic" serve as labels which provide a simple way to refer to a more complex behavioural profile. As such, they are not meant in the everyday sense of dispositions of behaviour or attitude toward life; rather they are meant to communicate which type of comparison (best case or worst case, as will be discussed in upcoming sections) the focal DM is effecting when applying initial state stabilities. Given that within the initial state stability framework it is necessary to compare states within sets, transitivity of preferences for each DM must hold.

Furthermore, similar to the distinction drawn between GMR and SEQ stabilities with respect to opponent moves, the opponent move sets considered can be sets of UMs or UIs. This consideration further divides the types of initial state stabilities into a total of four distinct solution concepts.

3.2.1 Pessimistic Initial State Stabilities

These solution concepts are meant to capture a risk-averse DM by comparing the worst-case scenario across outcome sets. In other words, for a starting state s_k, the DM will compare the least preferred state in the set of opponent moves from s_k to the least preferred states in the sets of opponent moves from the UIs available from s_k. If the least preferred state in s_k is equally or more preferred to the least preferred states in the other sets, then DM i will stay at s_k.

Formally, the set of least preferred states is defined as follows:

Definition 3.1. The set $w(s_t)$ denotes the set of least preferred states for DM i in the set $R_{N-i}(s_t)$. If $R_{N-i}(s_t) = \emptyset$, then $w(s_t) = s_t$.

Definition 3.2. A state s_k is *pessimistic initial state* (PIS) stable for DM i if $w(s_t) \precsim_i w(s_k)$ for all $s_t \in R_i^+(s_k)$.

This DM is verifying whether the worst-case scenario that could arise from staying at initial state s_k is more or equally preferred to the worst-case scenario arising from moving away from s_k by following an available UI. If the worst-case scenario at s_k is more preferred to all of the other worst-case scenarios, then DM i stays at the initial state s_k.

Analogously, if DM i is aware of their opponents' preferences, sets of opponent UIs can be considered:

Definition 3.3. The set $w^+(s_t)$ denotes the set of least preferred states for DM i in the set $R^+_{N-i}(s_t)$. If $R^+_{N-i}(s_t) = \emptyset$, then $w^+(s_t) = s_t$.

With this additional definition, we can formally define sequentially pessimistic inital state stability:

Definition 3.4. A state s_k is *sequentially pessimistic initial state* (SPIS) stable for DM i if $w^+(s_t) \precsim_i w^+(s_k)$ for all $s_t \in R_i^+(s_k)$.

The Elmira conflict detailed in Section 3.1.2 will be used to illustrate these calculations. First, consider state s_1 for LG. If LG were to remain at state s_1, then UR could move to states s_3 and s_9, while the MoE can move to state s_2. From s_2 and s_3, the conflict can be moved to state s_4 by the LG and UR, respectively. Hence $R_{UR,MoE}(s_1) = \{s_2, s_3, s_4, s_9\}$. If, on the other hand, LG were to take its UI to state s_5, then UR can move to states s_9 and s_7, while the MoE can move to state s_6, and UR and MoE can, by moving in sequence, also reach state s_8. Hence $R_{UR,MoE}(s_5) = \{s_6, s_7, s_8, s_9\}$. Among the states in $R_{UR,MoE}(s_1)$, LG's least preferred state is s_9, and s_9 is also the least preferred state for LG in $R_{UR,MoE}(s_5)$. Since $w(s_1)_{LG} = s_9$ and $w(s_5)_{LG} = s_9$, then LG is better off not moving from state s_1, which is PIS stable for LG.

If only opponent UIs are considered, one has $R^+_{UR,MoE}(s_1) = \emptyset$ and $R^+_{UR,MoE}(s_5) = \emptyset$; in other words, neither opponent has a UI from either state s_1 or state s_5 and would remain there. Comparing the two states, $w(s_5)_{LG} = s_5$ $w(s_1)_{LG} = s_1$, but $s_5 \succ_{LG} s_1$, therefore state s_1 is not SPIS stable for LG since they are better off moving to state s_5 from state s_1 rather than remaining.

On the other hand, state s_2 is SPIS stable for LG: $R^+_{UR,MoE}(s_2) = \{s_4, s_9\}$ and $R^+_{UR,MoE}(s_6) = \{s_8, s_9\}$. The least preferred state in the first set is s_4, and the least preferred state in the second set is s_9. Since $w^+(s_2)_{LG} = s_4$, $w^+(s_6)_{LG} = s_9$, and $s_4 \succ_{LG} s_9$, LG would do better by remaining at state s_2, hence s_2 is SPIS stable for LG.

3.2.2 Optimistic Initial State Stabilities

In contrast to PIS and SPIS, DM i could, rather than comparing worst-case outcomes, compare best-case outcomes. This DM is more risk-seeking than its pessimistic counterpart; what matters is not the worst-case, but what might happen in the best-case scenario.

Analogously to the sets $w(s_t)$ and $w^+(s_t)$, we can also define the sets of most preferred outcomes as follows:

Definition 3.5. The set $b(s_t)$ denotes the set of most preferred states for DM i in the set $R_{N-i}(s_t)$. If $R_{N-i}(s_t) = \emptyset$, then $b(s_t) = s_t$

Definition 3.6. The set $b^+(s_t)$ denotes the set of most preferred states for DM i in the set $R^+_{N-i}(s_t)$. If $R^+_{N-i}(s_t) = \emptyset$, then $b^+(s_t) = s_t$

It is now possible to define initial state stabilities in the optimistic case:

Definition 3.7. A state s_k is *optimistic initial state* (OIS) stable for DM i if $b(s_t) \precsim_i b(s_k)$ for all $s_t \in R^+_i(s_k)$.

In words, DM i is verifying whether the best possible outcome in $R_{N-i}(s_k)$ is better or equal to the best possible outcomes in $R_{N-i}(s_t)$ for all UIs s_t from s_k. If this is the case, then DM i will remain at s_k with the hope that the conflict will move to the best possible state.

The analogue definition with opponent UIs instead of opponent UMs is:

Definition 3.8. A state s_k is *sequentially optimistic initial state* (SOIS) stable for DM i if $b^+(s_t) \precsim_i b^+(s_k)$ for all $s_t \in R^+_i(s_k)$.

Once again referring to the Elmira conflict as an example, consider state s_3 for UR, who has the option to stay or move to state s_9. To check OIS stability, one finds $R_{LG,MoE}(s_3) = \{s_7, s_4\}$ and $R_{LG,MoE}(s_9) = \{s_9\}$. The most preferred state in the first set is state s_9, which is also trivially the most preferred state in the second set. Since $b(s_3)_{UR} = s_4$ and $b(s_9)_{UR} = s_9$, state s_3 is OIS stable for UR. Checking the same state for SOIS stability, one has $R^+_{LG,MoE}(s_3) = \{s_7\}$ and $R^+_{LG,MoE}(s_9) = \{s_9\}$. Since $b^+(s_3) = s_7 \succ_{UR} s_9 = b^+(s_9)$, state s_3 is not SOIS stable for UR.

State s_3 is, however, SOIS stable for LG: $R^+_{UR,MoE}(s_3) = \{s_9\}$ and $R^+_{LG,MoE}(s_7) = \{s_9\}$. Since $b^+(s_3)_{LG} = s_9$ and $b^+(s_7)_{LG} = s_9$, state s_3 is SOIS stable for LG.

Table 3.2 summarizes the new solution concepts and describes some of their important characteristics. PIS and SPIS are more conservative solution concepts since they are motivated by a desire to mitigate the worst-case scenarios; OIS and SOIS, on the other hand, are more risk seeking as they seek to maximise the best-case secnarios.

Table 3.2: Initial state solution concepts

Solution Concept	Description	Foresight	Know-ledge of Prefer-ences	Risk
PIS	Worst case resulting from staying is better than worst case resulting from opponent UM responses to all UIs	Medium	Own	Conservative
SPIS	Worst case resulting from staying is better than worst case resulting from opponent UI responses to all UIs	Medium	All	Conservative
OIS	Best case resulting from staying is better than best case resulting from opponent UM responses to all UIs	Medium	Own	Some risk
SOIS	Best case resulting from staying is better than best case resulting from opponent UI responses to all UIs	Medium	All	Some risk

3.3 Relationships Among Stability Concepts

There exist some relationships among the different initial state stabilities defined in the previous sections. These relationships are described and formalized in the following propositions.

Proposition 3.1. PIS stability is equivalent to OIS stability if $|R_{N-i}(s_k)| = |R_{N-i}(s_t)| = 1$ for all $t \in R_i^+(s_k)$.

Proof. Suppose that $|R_{N-i}(s_k)| = |R_{N-i}(s_t)| = 1$ for all $t \in R_i^+(s_k)$. This means that there is only one state in each of the sets of counter-moves by opponents from s_k and from s_t $\forall s_t \in R_i^+(s_k)$. In other words, whatever state s_x is in $R_{N-i}(s_k)$, it is by default both the most and least preferred state in that set: $w(s_k) = b(s_k) = s_x$. This is also true for the state $s_y \in R_{N-i}(s_t)$, i.e., $w(s_t) = b(s_t) = s_y$ for all $s_t \in R_i^+(s_k)$. Since $w(s_k)$ (=$b(s_k)$) is compared to $w(s_t)$ (=$b(s_t)$) for all $t \in R_i^+(s_k)$, calculating PIS stability for s_k is identical to calculating OIS stability for that state. □

Proposition 3.2. SPIS stability is equivalent to SOIS stability if $|R^+_{N-i}(s_k)| = |R^+_{N-i}(s_t)| = 1$ for all $t \in R_i^+(s_k)$.

Proof. The proof identical to that of Proposition 3.1 shown above except that the sets $R_{N-i}(s_u)$ are replaced by the sets $R^+_{N-i}(s_u)$ for $u = k, t$. □

Proposition 3.3. PIS stability is equivalent to SPIS stability if $R_{N-i}(s_k) = R^+_{N-i}(s_k)$ and $R_{N-i}(s_t) = R^+_{N-i}(s_t)$ for all $t \in R_i^+(s_k)$.

Proof. Suppose that $R_{N-i}(s_k) = R^+_{N-i}(s_k)$ and $R_{N-i}(s_t) = R^+_{N-i}(s_t)$ for all $t \in R^+_i(s_k)$. This means that from state s_k, all opponent moves happen to be UIs and that the same occurs for state s_t for all $t \in R^+_i(s_k)$. Since $R_{N-i}(s_k) = R^+_{N-i}(s_k)$, this means that $w(s_k) = w^+(s_k)$; since $R_{N-i}(s_t) = R^+_{N-i}(s_t)$ for all $t \in R^+_i(s_k)$, this means that $w(s_t) = w^+(s_t)\ \forall s_t \in R^+_i(s_k)$. Therefore, verifying PIS stability, which compares $w(s_k)$ to $w(s_t)\ \forall s_t \in R^+_i(s_k)$ is equivalent to SPIS stability, which compares $w^+(s_k)$ to $w^+(s_t) \forall s_t \in R^+_i(s_k)$. □

Proposition 3.4. OIS stability is equivalent to SOIS stability if $R_{N-i}(s_k) = R^+_{N-i}(s_k)$ and $R_{N-i}(s_t) = R^+_{N-i}(s_t)$ for all $t \in R^+_i(s_k)$.

Proof. The proof is identical to that of Proposition 3.3 except that the sets $w(s_u)$ and $w^+(s_u)$ are replaced with the sets $b(s_u)$ and $b^+(s_u)$ for $u = t, k$. □

As shown in Propositions 3.1 to 3.4, there are special cases in which the new stability concepts can be reduced to one another. PIS (SPIS) is equivalent to OIS (SOIS) if there is only one state in each of the sets of opponent counter-moves; by default those states will be both the least and most preferred states in the set. PIS (OIS) is equivalent to SPIS (SOIS) if all opponent moves happen to be UIs. These results are useful for calculations since these can, in some cases, be greatly simplified.

Proposition 3.5. If a state $s_k \in S$ is Nash stable for DM i, then it is also

(a) PIS;

(b) SPIS;

(c) OIS; and

(d) SOIS stable for DM i.

Proof. (a) : Suppose that s_k is Nash stable for DM i, then $R^+_i(s_k) = \emptyset$. This means that the opponent sets $R_{N-i}(s_t)$ are also empty, since there is no $s_t \in R^+_i(s_k)$. The condition that $w(s_t) \precsim_i w(s_k)$ for all $s_t \in R^+_i(s_k)$ is thus trivially satisfied.
(b) - (d): The remaining proofs are analogous to the one shown above. The key point is that a Nash stable state has no UIs and that this trivially satisfies the definitions of SPIS, OIS, and SOIS stabilities. □

Proposition 3.6. If opponents have no UMs from s_k and $R_{N-i}(s_t) \neq \emptyset$ for all $s_t \in R^+_i(s_k)$, then PIS stability is equivalent to GMR stability.

Proof. The calculation of PIS stability compares $w(s_k)$ to $w(s_t)$ for all $s_t \in R^+_i(s_k)$. If opponents have no UMs from s_k, the comparison will simply be between $w(s_k) = s_k$ and $w(s_t)\ \forall s_t \in R^+_i(s_k)$. That is, the initial state s_k is compared to states in $R_{N-i}(s_t)$ for $s_t \in R^+_i(s_k)$.

Note that this comparison is identical to that being performed in GMR calculations: state s_k is GMR stable if for every $s_t \in R^+_i(s_k)$, there is at least one state in $R_{N-i}(s_t)$ which is less preferred to s_k. That is, at least one state in $R_{N-i}(s_t)$ is less preferred to s_k for every $s_t \in R^+_i(s_k)$.

Having established that PIS and GMR stabilities compare the same sets, it remains to show that the comparisons will yield the same stability outcomes. State s_k is PIS stable if $w(s_t) \precsim_i w(s_k) = s_k\ \forall s_t \in R^+_i(s_k)$. Another way to express this is that s_k is less preferred to at least one state (call it $w(s_t)$) for each of the sets $R_{N-i}(s_t)$; this is precisely the definition of GMR stability. □

Proposition 3.7. If opponents have no UIs from s_k and $R_{N-i}(s_t) \neq \emptyset$ for all $s_t \in R_i^+(s_k)$ then SPIS stability is equivalent to SEQ stability.

Proof. The proof is analogous to that of Proposition 3.6 with $R_{N-i}(s_t)$ replaced with $R_{N-i}^+(s_t)$ and $w(s_u)$ replaced with $w^+(s_u)$ for $u = t, k$. □

Propositions 3.5 to 3.7 relate the new solution concepts to the core Graph Model stability concepts. These results are useful for calculating initial state stabilities: for example, if that a state is Nash stable, one immediately knows that it is also PIS, SPIS, OIS, and SOIS stable. Furthermore, under the conditions specified in Propositions 3.6 and 3.7, it is possible to reduce PIS and SPIS calculation to GMR and SEQ calculations, respectively.

3.4 Application: Elmira conflict

In this section, the complete stability results for the conflict described in Section 3.1.2 are detailed. Both initial state and standard Graph Model stabilities are calculated and discussed.

3.4.1 Initial State Stabilities

The calculations for rest of the Elmira conflict are carried out in a manner similar to the one detailed in Sections 3.2.1 and 3.2.2. Table 3.3 summarizes the equilibrium results by type of equilibrium. The detailed calculation results are shown in Table 3.4 which is included at the end of this section for reference.

Table 3.3: Initial state stability equilibria for the Elmira conflict

Type of equilibrium	States
PIS	s_1, s_4, s_5, s_8, s_9
SPIS	s_3, s_5, s_8, s_9
OIS	s_4, s_5, s_7, s_8, s_9
SOIS	s_5, s_8, s_9

Given the DMs' individual stability results, there are PIS equilibria at states s_1 (UR delays the appeal), s_4 (MoE modifies the CO and UR accepts it), s_5 (UR appeals and LG insists), s_8 (MoE modifies the CO, UR accepts it, and LG insist), and s_9 (UR abandons the site). In other words, if all three DMs were risk-averse and concerned about opponent sanctioning, one of these states would be expected to be the outcome of the conflict.

If the DMs were risk-averse and also aware of each others' preferences, then states s_3 (UR accepts CO), s_5, s_8, and s_9 are SPIS equilibria. Note that states s_5, s_8 and s_9 are both PIS and SPIS equilibria; this means that the DMs' knowledge of their opponent's preferences does not affect the stability of these states. The opposite is true for states s_1, s_4, and s_3: states s_1 and s_4 are equilibria only under the assumption that the possibility of opponent sanctioning by UMs is allowed, while s_3 is an equilibrium only when opponent UIs are considered.

Looking at the OIS equilibria, these occur at states s_4, s_5, s_7 (LG insists and UR accepts CO), s_8 and s_9. If all three DMs were risk-taking and concerned about opponent UM sanctioning, those states represent possible resolutions of the conflict.

The SOIS equilibria for this conflict occur at states s_5, s_8, and s_9. Note that each of these states is included in the list of OIS equilibria; for these states, assumptions regarding knowledge of opponent preferences do not affect the stability of the states. It thus remains that states s_4 and s_7 are affected by whether opponent preferences are known.

Now, comparing PIS and OIS equilibria, note that states s_4, s_5, s_8, and s_9 appear in both lists. This indicates that the level of optimism/pessimism for the DMs does not affect these states as they are equilibria either way. State s_1 is exclusively a PIS equilibrium, while state s_7 is purely an OIS equilibrium; the likelihood of these states being the resolution of the conflict thus depends on the likelihood that each of the DMs follows either PIS or OIS reasoning.

A similar comparison of SPIS and SOIS equilibria shows that they have states s_5, s_8, and s_9 in common. Once again, this shows that the DMs' behaviour is not affected by whether they are optimistic or pessimistic. State s_3 is exclusively a SPIS equilibrium: for it to be the outcome of the conflict, all DMs must follow SPIS reasoning.

Remark that states s_5, s_8, and s_9 are equilibria under all initial state stability concepts. This is not particularly surprising for state s_9 since a move to this state is irreversible. For s_5 and s_8, however, this is important to note since it is thought that states which are equilibria under a range of solution concepts are more likely to occur in practice (Kilgour, Hipel, Fang and Peng 2001; Kilgour and Eden 2010). The initial state stability analysis thus points to these states as the most likely outcomes. States s_4 (PIS and OIS equilibrium), s_1 (PIS equilibrium), s_3 (SPIS equilibrium), and s_7 (OIS equilibrium) are also possiblities.

It may also be useful to examine the results from each individual DM's point of view. If each DM's two most and two least preferred states are examined, it is interesting to see what kind of advice each type of initial state stability provides them with. The MoE's two most preferred states are s_7 and s_3, while their two least preferred states are s_9, and s_6. In all cases, these states are PIS, SPIS, OIS, and SOIS stable; this means that the MoE should remain at these states when they occur, regardless of their own level of risk taking.

UR's two most preferred states, s_1 and s_4 are also stable under all initial state stabilities, meaning that UR should remain at these states if they should occur. As for UR's two least preferred states s_7 and s_6, state s_6 is not stable under any initial state stability concept, meaning that UR should move away from s_6 regardless of its level of risk aversion. State s_7, however is stable under OIS; this means that UR should stay at state s_7 only if their decision-making follows OIS.

Finally, the stability results show that LG's most preferred state s_7 is stable under all initial state stability concepts; however, state s_3, its second most preferred state, is OIS unstable. This means that if LG is optimistic and not concerned about opponent sanctioning, it should move from s_3. LG's least preferred state, s_9 is stable under all initial state stability concepts; however, its second least preferred state s_2 is only PIS and SPIS stable. This means that if LG is risk taking, it should move away from s_2.

Table 3.4: Initial state stabilities for the Elmira conflict

	PIS			SPIS			OIS			SOIS		
	MoE	UR	LG	MoE	UR	LG	MoE	UR	LG	MoE	UR	LG
s_1	✓	✓	✓	✓	✓	x	✓	✓	x	✓	✓	x
s_2	✓	x	✓	✓	x	✓	✓	x	x	✓	x	x
s_3	✓	x	✓	✓	✓	✓	✓	x	x	✓	x	✓
s_4	✓	✓	✓	✓	✓	x	✓	✓	✓	✓	✓	x
s_5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
s_6	✓	x	✓	✓	x	✓	✓	x	✓	✓	x	✓
s_7	✓	x	✓	✓	x	✓	✓	✓	✓	✓	x	✓
s_8	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
s_9	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

A checkmark (✓) indicates that the given state is stable under the specified solution concept, while an "x" means that the state is not stable for that solution concept.

3.4.2 Comparison to Graph Model Stabilities

The complete stability results of the Graph Model for the Elmira conflict are shown in Table 3.5. As can be seen, states s_5, s_8, and s_9 are Nash equilibria, states s_5, s_8, and s_9 are SEQ equilibria, states s_1, s_4, s_5, s_8, and s_9 are GMR equilibria, and states s_1, s_4, s_5, s_8, and s_9 are SMR equilibria.

Table 3.5: Graph Model stabilities for the Elmira conflict

	Nash			SEQ			GMR			SMR		
	MoE	UR	LG	MoE	UR	LG	MoE	UR	LG	MoE	UR	LG
s_1	✓	✓	x	✓	✓	x	✓	✓	✓	✓	✓	✓
s_2	✓	x	x	✓	x	✓	✓	x	✓	✓	x	✓
s_3	✓	x	x	✓	x	✓	✓	x	✓	✓	x	✓
s_4	✓	✓	x	✓	✓	x	✓	✓	✓	✓	✓	✓
s_5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
s_6	✓	x	✓	✓	x	✓	✓	x	✓	✓	x	✓
s_7	✓	x	✓	✓	x	✓	✓	x	✓	✓	x	✓
s_8	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
s_9	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

By comparing the initial state stability results with those from the Graph Model, one can see that, in addition to being PIS, SPIS, OIS, and SOIS equilibria, states s_5, s_8, and s_9 are also Nash, SEQ, GMR, and SMR equilibria. This furthers the intuition that one of these states is likely to be the resolution of the conflict. In fact, the conflict historically ended in state s_8 where a favourably modified CO was negotiated by the MoE and UR.

It should also be noted that using the initial state stability framework revealed equilibria that were not predicted by any of the Graph Model solution concepts, namely state s_3 (SPIS equilibrium) and state s_7 (OIS equilibrium). This type of information is useful to analysts since it avoids potential surprises that stem from only considering standard Graph Model solution concepts. Armed with this knowledge, analysts can determine whether state s_3 or s_7 is likely to occur given what is known about the DMs.

The example thus shows that new equilibria can be discovered by using initial state stability solution concepts. It is well known that analysing a conflict using a range of solution concepts allows one to better identify likely outcomes (Madani 2013). Initial state stabilities thus contribute to analysts' behavioural insights in determining possible conflict resolutions.

3.5 Chapter Summary

A new family of stability concepts, initial state stabilities, is introduced to complement existing Graph Model stability concepts. This addresses the first of the two motivating questions discussed in the introduction: "What are other ways to define stability for a state?". The new stability is defined by introducing a broader state comparison technique: comparisons are executed between the possible consequences of leaving a given starting state and of staying at that state. Pessimistic initial state stabilities compare worst case scenarios; a state is stable if the worst case scenario that it generates is more preferred than the worst case scenarios that occur at each of that DM's UIs. Optimistic initial state stabilites, on the other hand, compare best case scenarios: a state is stable if the best case scenario that could result from it is more preferred to the best case scenarios which could occur at each of the DM's UIs. For both types of solution concepts, opponent UMs (PIS, OIS) and UIs (SPIS and SOIS) are considered, giving a total of four new solution concepts.

The definition of these new and useful solution concept contributes to both the literature and, more broadly, to decision-makers' tool-kits. Building upon solution concepts based on the Graph Model strengthens the analysis and insights garnered from this framework; these can then be passed on to decision-makers. Furthermore, the practice of analysing conflicts using a range of solution concepts can result in a better understanding of the conflict and of its possible resolutions.

By using initial state stability to analyse a dispute, DMs can determine whether staying at their most preferred state could result in a worse outcome in the future or whether moving away from their least preferred state could yield a better outcome in the future. These new insights allow DMs to answer more than the usual question of "Should I stay because moving might be risky?"; they can now ask themselves questions such as "If I stay, might I be sanctioned for it later on?" and "If I go, might I end up in a better state later on?". These questions provide DMs with new avenues to explore and to analyse when thinking strategically.

As demonstrated by the Elmira conflict case study, using initial state stabilities to analyse a real-world dispute is fairly straightforward and, for small conflicts, can be done by hand. Having the conflict represented in Option Form and in Graph Model form greatly facilitates these calculations. The results gained from the analysis are useful to DMs and analysts: gaining additional stability information about an individual state is crucial to providing improved strategic insights about the possible evolution of the conflict.

Chapter 4

Inverse Engineering

This chapter focuses on the second motivating question posed in Chapter 1, in which knowledge of a DM's preferences is not assumed. This assumption is quite reasonable: researchers who work with the Graph Model know that preferences are notoriously difficult to obtain. The discovery of each DM's preferences is important for analysing both historical and ongoing conflicts as the analytical results depend on their accuracy. For ongoing conflicts, a DM may not be willing to share their preferences with an analyst and will be much less willing to do so with an opponent; the study of historical conflicts, on the other hand, may be hampered by a lack of information or reliable sources. There is also the possibility that a DM may be uncertain about their own preferences, taking actions but perhaps not clearly understanding why.

Preference information is crucial to DMs and analysts who wish to garner useful and informative insights about a conflict. Although preferences can be difficult to ascertain, opponent moves are clear and observable. Since moves are informed by a DM's preferences, the transparency of an opponent's moves can be leveraged by a DM to infer that opponent's preferences. In this chapter, methods for ascertaining DM preferences using observed behaviour are proposed based on different DM profiles. The methods narrow down the set of a DM's possible preference rankings using the information provided by that DM's actions. The results obtained from these techniques are to produce more accurate insights when using the Graph Model. Furthermore, a clear read on opponent preferences allows for improved decision-making; one will have a better idea of how to arrive at a more preferable result either by changing their opponents' preferences or by moving strategically.

Adopting an inverse approach for preference elicitation within the Graph Model is a new endeavour. Such a procedure has several advantages: first, there is little to no need for direct consultation between DMs and analysts since a DM's preferences can be inferred from their behaviour alone. This is particularly relevant for adversarial situations in which DMs may not wish to communicate their preferences to analysts in order to preserve a strategic advantage. This characteristic is also helpful in the study of historical conflicts for which documentation might be sparse. Rather than relying on an analyst's best guess or on an approximation of DM preferences, empirical data are used to generate insights.

Second, the dynamic nature of the algorithms allows for ongoing analysis which adapts to the latest changes in the conflict. This contrasts with the static approaches characteristic of the Graph Model which typically select a single point in time to analyse.

Third and most important, the inverse approach can be used to provide relevant advice to DMs about what to do next. The ADVICE function which is the focus of Section 4.5 is designed to provide key information such as expected value and probability of occurrence. The overarching goal is to assist with real-time decision-making: advice is updated as the conflict progresses, and DMs can use the information provided by opponent moves to determine their best countermoves and strategies. The goal is not to replace a DM in the decision-making process, but to supply enriching information for reaching more informed decisions.

This chapter is organized as follows: Section 4.1 discusses the philosophical implications of relating preferences and options; Section 4.2 provides an overview of the literature and work done on preference elicitation; Sections 4.3, 4.4, and 4.5 detail the methodology and its associated algorithms; and Section 4.6 details a simple case study.

4.1 Philosophical Interlude

In this section, the fundamentals of Revealed Preference Theory are discussed and then critiqued. In particular, the classical economics view of preferences is contrasted with the perspective provided by psychology and behavioural economics. While the former strongly informs the methodology introduced in the upcoming sections, the latter offer valid criticisms which deserve to be presented. The synthesis section provides a way forward which acknowledges both the benefits and shortcomings of the economics approach to preferences.

4.1.1 Revealed Preference Theory

From a classical economics standpoint, mental preferences describe the attitude of a DM toward a set of objects. A DM's choices always maximise the utility of their behavioural preference relation i.e., the preference relation which is observable to the analyst. A DM's actions manifest their behavioural preferences, which are assumed to be identical to their mental preferences.

The latter statement is the cornerstone of the standard economic or neo-classical economic approach in which mental preferences and behavioural preferences are in perfect correspondence. This Principle of Revealed Preference was originally devised to infer consumer behaviour based on their choice of goods and posits that observed behaviour can reveal a DM's mental preferences (Samuelson 1938; Samuelson 1948).

A corollary of the Principle is the Weak Axiom of Revealed Preference which ensures that a DM is consistent in their preferences: if a DM chooses option A over option B, they must not then choose option B over option A since their first choice revealed a preference for option A over option B. In this view, actions are dictated by preferences; DMs make choices according to whether or not they maximise the utility of their mental preferences.

Although Revealed Preference Theory has been further developed and enriched since Samuelson's day (see Hands 2013 for a history and overview), the Principle of Revealed Preferences is what underlies the methodology and algorithms presented in the upcoming sections; for conciseness, additional details and subtleties of Revealed Preference Theory will not be discussed here. The key assumptions to keep in mind are (1) DMs are rational (i.e., are utility maximisers) and (2) mental preferences are identical to behavioural preferences.

In the context of the Graph Model, DMs are choosing their actions among their set of options (not among a bundle of goods as in Samuelson's original formulation). Preferences are understood as rankings of a set of outcomes for a given conflict: each DM has their own individual set of mental preferences which describe their attitude toward the set of states. During a conflict, DMs must make choices about which of their options to take (or not). The choices made by DMs (i.e., their behavioural preferences) can directly reveal their mental preferences.

4.1.2 Critiques of Revealed Preference Theory

Research in psychology and in behavioural economics has demonstrated that the relationship between actions and preferences is not as straightforward as classical economists claim. Several empirical studies have demonstrated an inverse causality to the one outlined in the previous section: actions can, in fact, affect preferences (see, for example: Festinger and Carlsmith 1959, Ariely and Norton 2008, Sharot, Martino and Dolan 2009, and Egan, Bloom and Santos 1959).

From the economics side, Revealed Preference Theory has been criticised by, amongst others, Sen and Hausman (Sen 1973; Sen 1993; Hausman 2000). One of the arguments that they formulate against Revealed Preference Theory is that choice alone is not sufficient to reveal preferences. Sen's argument centres around the violation of the axioms of revealed preference that occur when contexts change. Sen terms these changes of context menu dependence (Sen 1993). Since the DM's choices often depend on the context, argues Sen, they alone cannot determine that DM's preferences.

Hausman's critique follows a similar line: a DM's choice is not enough to reveal their preference as there may be other influencing factors in play. For instance, the DM may make mistakes, may have an incomplete ranking (i.e., some elements may not be comparable to one another), or may have moral considerations to take into account (Hausman 2000).

To summarize the important criticisms highlighted here, first we note that the relationship between actions and preferences is not unidirectional, but can often form a closed feedback loop; and second, a DM's choice may not be sufficient to reveal their preference due to context or other mitigating factors – in other words, DMs may behave in ways inconsistent with pure utility maximisation.

4.1.3 Synthesis

Recall the two assumptions highlighted in Section 4.1.1: (1) DMs are rational (i.e., are utility maximisers) and (2) mental preferences are identical to behavioural preferences. In light of the criticisms discussed in Section 4.1.2 both of these must be carefully re-examined.

The first assumption regarding utility maximisation requires relaxing; as Sen and Hausman correctly point out, DMs may take actions on the basis of some other heuristic. This is also highlighted in Rubinstein and Salant's appraisal of Revealed Preference Theory which argues that the theory is useful insofar as it helps to classify decision models which incorporate additional psychological factors (Rubinstein and Salant 2008). Fortunately, a variety of heuristics is incorporated into the new methodology, which introduces behavioural profiles (see section 4.4) such as GMR and SMR which do not always assume that the DM chooses the state which maximizes their payoff. As will be shown, these alternative profiles allow a DM to choose actions based on anticipated opponent counter-moves and sanctions rather than on the basis of simple payoff comparison. Furthermore, the incorporation of further decision heuristics remains an active part of this research.

What psychology and behavioural economics reveal about the relationship between preferences and actions, namely that they have the potential to be mutually influencing, can make it tempting to discard any results obtained with assumption (2) in mind. However, as Spiegler argues, Revealed Preference Theory is a useful model development tool: "Even if one rejects the revealed preference principle as a criterion for determining the admissibility of “behavioral” decision models, the principle still has a heuristic value in the development of such models. A rudimentary revealed-preference exercise helps clarifying general aspects of the behavior induced by the model. The clarifcation obtained in this way is so basic, that it cannot be left for a future decision theorist. Instead, it should be part of the behavioral theorist’s bag of tools" (Spiegler 2008). Thus, taking the results obtained from a Revealed Preference Theory with a grain of salt while keeping present their heuristic value offers a way forward.

The goal of this section is to better situate the methodology which will be introduced later in this chapter both in the context of Revealed Preference Theory and of its associated criticisms. As discussed, the methodology allows for DM behaviour that falls outside of the traditional utility maximisation heuristic by introducing behavioural profiles such as GMR and SMR. Furthermore, although the relationship between actions and preferences is indeed more complex than this method assumes, it nevertheless provides a benchmark for analysis and can be one of the many tools used by analysts to gain a better understanding of the conflict.

4.2 Review of the Literature

Preference elicitation, that is, determining a DM’s preferences over a list of objects or outcomes, is a key step in constructing decision-making and game-theoretic models. As a result, both "forward" and "inverse" approaches have been developed to assist analysts and DMs with this important task. Generally speaking, the former attempt to construct preference rankings based on attributes, values, or judgments provided by the DM; common methods in this vein include the analytic hierarchy process (AHP) (R. W. Saaty 1987; T. L. Saaty 1980), ELimination and Choice Expressing the REality (ELECTRE) (Figueira et al. 2013; Roy 1968) and value-focused thinking (Keeney 1992).

Inverse approaches, on the other hand, work "backwards" from observed behaviours to a reward function which contains preference information. These approaches commonly use Markov Decision Processes (MDPs) or Partially Observable MDPs (POMDPs) to infer a DM’s reward function given an observed behaviour (Abbeel and Ng 2004; Ng and Russell 2000; Boutilier 2002). Such techniques are designed to solve problems in inverse reinforcement learning: one can only observe behaviours from an expert and must use these observations to learn to perform these same tasks.

Preference elicitation is a critical step in the Graph Model methodology; the model relies on preference input from all of the DMs in order to produce its analysis results. Within the Graph Model paradigm, "forward" preference elicitation is akin to eliciting ordinal preferences from DMs. Several techniques can be used to perform this task, including option weighting, option prioritizing, and direct ranking (Fang et al. 2003b). More recently, work has been done to integrate AHP within the Graph Model (Ke, Li and Hipel 2007; Ke, Li and Hipel 2012; Ke, Fu et al. 2012) and to extend option prioritization to non-standard preference structures (Bashar, Kilgour and Hipel 2014). In this case, the analyst is working closely with all of the DMs involved in the conflict and, in doing so, is able to elicit the necessary preference information to carry out a stability analysis.

Many techniques have examined preferences and their role in the Graph Model. One family of procedures consists of exploring non-standard preference relations in order to capture a DM's uncertainty over the set of states. Such apporaches include unknown preferences (Li, Hipel et al. 2004; Xu, Hipel, Kilgour and Chen 2010), fuzzy preferences (Bashar, Hipel and Kilgour 2012), probabilistic preferences (Rego and dos Santos 2015), grey-based preferences (Kuang et al. 2015), strength of preferences (Hamouda, Kilgour and Hipel 2004; Hamouda, Kilgour and Hipel 2006; Xu, Hipel and Kilgour 2009b), and information-gap models (Ben-Haim and Hipel 2002).

The first four of these approaches examine different preference relations and structures in order to deal with the uncertainty that a particular DM has over the set of states. Strength of preferences considers "strongly preferred" relations between states in addition to the usual "preferred" and "indifferent" relations (Hamouda, Kilgour and Hipel 2004; Hamouda, Kilgour and Hipel 2006); this concept was later generalized to any level of preference (Xu, Hipel and Kilgour 2009b) and incorporated into coalition analysis (Li, Inohara and Xu 2014).

Unknown preferences make use of an additional "unknown" relation between possible states. Unknown preferences and strength of preference are also combined to produce a hybrid preference structure (Xu, Hipel, Kilgour and Chen 2010). Fuzzy preferences make use of fuzzy sets, which specify a degree of belonging of an element to a set, to define fuzzy relations and fuzzy preferences (Bashar, Hipel and Kilgour 2012). In the probabilistic Graph Model, it is assumed that DMs have probabilistic preferences over the set of states, whereby one state is preferred over another with a certain probability (Rego and dos Santos 2015). Grey-based preferences use grey numbers, which are real numbers that may be members of a discrete set of real numbers, or may fall within one or several intervals (Kuang et al. 2015). In all cases, the new preferences structures are used to construct their respective analogues of Graph Model solution concepts and equilibria. The common assumption underlying most of these methods is that the DMs are uncertain of their or their opponents' ranking of states; some states may be strictly preferred over others, but some ambiguity may exist. In other words, the uncertainty may be present within each DM's own preference ranking.

Another family of approaches attempts to bypass preferences entirely. Incomplete information in the Graph Model is examined by Sakakibara et al (Sakakibara, Okada and Nakase 2002). This approach specifies the minimum information needed for analysts or third parties to perform stability analyses. The robustness analysis developed in this work is general enough to be applied to conflicts in which DM preferences are not fully known. This work's primary goal is not to discover DM preferences, but to allow for stability analysis of conflicts with incomplete information.

Yet a different method, the inverse Graph Model, deals with the problem of preference elicitation by identifying preference rankings that could lead to a desired resolution(Kinsara, Kilgour and Hipel 2014; Kinsara, Petersons et al. 2015; Wang et al. 2018). Rather than working forward from DMs, options, and preferences to conflict resolutions, this approach begins by selecting a desired resolution and determining which preferences are required to bring it about. Analysts use what information they have about the DMs' preferences and then apply the inverse approach to determine what other preference rankings are necessary for creating the desired resolution. Rather than attempting to ascertain opponent preferences based on actions, this methodology bypasses the problem entirely.

The philosophy guiding the present research falls somewhere in between the two families of approaches; no new preference structures are used, nor is the existing preference structure modified, yet the determination of preferences is not bypassed – it is the central concern. The key question is how to use observable behaviours to infer preferences, and how to use that information to the benefit of DMs. In order to answer this question, a simple preference structure is assumed and coupled with the knowledge gained from observing actions.

Adopting an inverse approach for preference elicitation in the Graph Model is a new endeavour (Garcia and Hipel 2017; Garcia, Obeidi and Hipel 2018). This procedure has several advantages: first, there is a lesser need for direct consultation between DMs and analysts since a DM's preferences and some useful accompanying information can be inferred from behaviour alone. This is particularly relevant for situations in which DMs wishing to preserve a strategic advantage may not wish to communicate their preferences to analysts. This is also helpful in the study of historical conflicts for which documentation might be sparse. Rather than relying on an analyst's best guess or on an approximation of DM preferences, empirical data are used to generate insights.

Second, the dynamic nature of the algorithms allows for ongoing analysis which adapts to the latest changes in the conflict. This contrasts with the static approaches of the Graph Model which typically selects a single point in time for analysis. Finally, the inverse approach can be used to provide relevant advice to DMs about what to do next. The Advice function which is the focus of Section 4.5 is designed to provide key information and is updated as the conflict progresses. DMs can use the information provided by observable moves to determine their best counter-moves and strategies in real time. The goal is not to replace a DM in the decision-making process, but rather to supply enriching information for reaching more informed decisions.

4.3 Details on Preference Rankings

The first part of this section emphasizes preference rankings, which are fundamental to the algorithms developed in the latter part. Then, once the problem is discussed in general terms, algorithms specific to each of three Graph Model stability concept are developed, consisting of Nash (Section 4.4.2), GMR (Section 4.4.3), and SMR (Section 4.4.5) stability.

Given the importance of preference rankings to the approach, it is valuable to delve into them a bit deeper. Preference rankings are weak orderings of the set of states, with states being equally preferred ($\sim$) or more preferred ($\succ$) to one another. For example, a DM might have the following preference ranking over a set of four states: $s_1 \succ s_2 \sim s_3 \succ s_4$. This means that state s_1 is preferred to state s_2, which is equally preferred to state s_3 and more preferred to state s_4.

Other jargon is employed to label preference rankings in the game theory and decision-making literature. For instance, ordinal preference refers to a weak ordering of states in which states are ordered or ranked from most to least preferred in which ties are allowed. Strict ordinal preference indicates an ordering of states which are ranked from most to least preferred without any ties being present. Both ordinal (weak ordering) and strict ordinal (ordering) of states satisfy the property of transitivity: $s_1 \lesssim s_2$ and $s_2 \lesssim s_3$ implies $s_1 \lesssim s_3$ for all $s_1, s_2, s_3 \in S$.

Preference rankings over the set of states can be mapped to payoff values: a DM's preference ranking can be recovered by comparing the payoff values of states and ordering them from highest payoff (i.e., most preferred) to lowest payoff (i.e., least preferred). The payoff that DM i receives from some state $s \in S$ is denoted by $\pi_i(s)$. Payoffs can, in theory, be any real number as long as the ranking is respected. To illustrate, for DM A, a valid payoff function for preference ranking $p_A = s_1 \succ s_2 \sim s_3 \succ s_4$ is $\pi_A(s_1) = 4$, $\pi_A(s_2) = \pi_A(s_3) = 2.5$, and $\pi_A(s_4) = 0$. A non-valid payoff function is $\pi_A(s_1) = 4$, $\pi_A(s_2) = 3$, $\pi_A(s_3) = 2$, and $\pi_A(s_4) = 2.5$ because $\pi_A(s_2)$ and $\pi_A(s_3)$ must be equal and since $\pi_A(s_3)$ must be greater than $\pi_A(s_4)$.

It is important to remember that preference rankings do not specify the degree of preference, only the preference order – it is not known how much more one state is preferred to another. Given this observation, the choice of payoff values is somewhat arbitrary. For the present purposes, an additional requirement is imposed: any payoff function taking values of $k_0\epsilon < k_1\epsilon < k_2\epsilon < \ldots < k_s\epsilon$ for $\epsilon > 0$ will do as long as there is equal distance between each payoff value, i.e., $|k_i - k_{i+1}| = |k_{i-1} - k_i| \ \forall i$. For simplicity, payoffs are set using $\epsilon = 1$ and $|k_i - k_{i+1}| = |k_{i-1} - k_i| = 1$. For example, the payoffs for DM A in the ranking given above are: $\pi_A(s_1) = 4$, $\pi_A(s_3) = \pi_A(s_2) = 3$, and $\pi_A(s_1) = 2$. Given this convention, states in the same equivalence class will receive the same payoff and the payoff of any given state is the total number of states minus that state's equivalence class index which is counted starting from the closest equivalence classes to the most preferred state.

Since valid preference rankings are weak orderings over the set of possible states, the total number of preference rankings is given by the number of partitions of the set of states into different equivalence classes (delimited by $\sim$) multiplied by the number of ways to order the elements within them. The total number of preference rankings of m states, allowing equalities, is given by the ordered Bell number (Good 1975):

$$Bell(m) = \sum_{k=0}^{s} k! \left\{ {m \atop k} \right\}, \tag{4.1}$$

where $\left\{ {m \atop k} \right\}$ is the Stirling number of the second kind which calculates the number of ways to partition a set of m objects into k non-empty subsets:

$$\left\{ {m \atop k} \right\} = \frac{1}{k!} \sum_{j=0}^{k} (-1)^{k-j} \binom{k}{j} j^m. \tag{4.2}$$

Thus, given the number of states in the conflict under study, it is possible to generate the list of all possible preference rankings. Although the list is initially quite large, observing a DM's moves allows rankings to be eliminated, shortenting the list of possible preference rankings. Each observed action has the potential to provide additional information about that DM's preference ranking.

The algorithms proceed by narrowing the set of feasible preference rankings at each iteration. The set $\mathcal{P}_i^{(k)}$ contains all possible preference rankings over m states for DM i after k iterations of the algorithm. If all preference rankings are assumed to be feasible at the start of the game, then $\mathcal{P}_i^{(0)} = \mathcal{P}^{(0)} \ \forall i$, where $\mathcal{P}^{(0)}$ is the set of all possible preference rankings over the set of states. As the conflict progresses, the cardinality of $\mathcal{P}_i^{(k)}$ decreases.

4.4 Behavioural Profiles

Let p_i denote DM i's preference ranking over the set of states and g^i denote i's action which is observed by the other DM(s). DM i's choice of action can convey information to the other DMs depending on i's behaviour profile. Behavioural profiles simply refer to the different types of DMs one might encounter. The behavioural profiles discussed in this paper correspond to the Graph Model stability concepts outlined in the previous section: for example, a DM may move according to Nash stability. Thus, DM i has chosen action g^i because it meets the criteria set by its behavioural profile. Therefore, for a behavioural profile B and some result R, DM i chooses g^i such that $R(g^i|p_i, B_i)$ is optimized. Knowing this, the remaining DMs can surmise that the preference rankings in which action g^i does not optimize the desired result (which is itself based on the behaviour profile) are removed from the initial set $\mathcal{P}_i^{(0)}$ of possible rankings to generate an updated list, $\mathcal{P}_i^{(1)}$.

Given the previous observation regarding the ability to map preference rankings to payoffs, the algorithms work directly with payoffs in order to facilitate comparisons and calculations. It is also assumed that the list of all possible payoffs has already been initialized and that the sets *unilateralMoves* (for all DMs) and *opponentUMs* can be generated.

The pseudocode for the RemovePayoffs function is shown in Algorithm 1. The foundation of this algorithm lies in the call of the isMoveSanctioned Boolean function, which is analysed in the next sections. The RemovePayoffs function essentially verifies for which possible rankings the observed action is consistent with the opponent's assigned behavioural profile. Rankings for which such an action is not compatible are removed from that opponent's list of possible rankings.

Algorithm 1 RemovePayoffs($startState$, $endState$)

```
for ranking ∈ possibleRankings do
    if isMoveSanctioned(startState, endState, ranking) then
        remove ranking
    end if
end for
```

4.4.1 Default Behavioural Profile

In the most general case, the opponent's behaviour profile can be completely unknown; a DM's reasoning for staying at or moving from a given starting state is opaque. In this case, a reasonable assumption is made about the opponents: they will never move to a state which is less preferred to the state that it is starting from. In other words, it is assumed that the opponent will not move in a way that might harm them.

Given a starting state s, there are two possibilities: either DM i stays at state s or DM i moves away from state s to state t. In the former case, no information is gained; i's motivation for staying at s is unknown. However, in the latter case, it can be conjectured that state t, is greater than or equally preferred to s. Rankings in which $t \prec s$ can thus be removed from the list of the DM's possible rankings. Algorithm 2 illustrates the procedure for removing payoffs when an opponent's profile is set to default.

Algorithm 2 isMoveSanctioned (Default)($startState$, $endState$, $ranking$)

if $startState > endState$ **then**
 return True
else
 return False
end if

A move is sanctioned for a default type DM if the end state is less preferred to the start state. Thus, when called by RemovePayoffs function, payoffs in which this condition is met are discarded: since the DM moved from $startState$ to $endState$, rankings in which $endState$ is strictly less preferred to $startState$ are excluded since it is assumed that the DM never moves to a less preferred state.

To summarize, the default profile behaves as follows:

- If DM i stays at s, no rankings are removed because no information is gained
- If DM i moves away from s to a state t, remove rankings in which $s \succ_i t$ because it must be that $t \succsim_i s$

4.4.2 Nash Profile

Recall that a state s is Nash stable if and only if there is no unilateral improvement from s. Since Nash stability is limited to the focal DM's moves, it does not require knowledge of the opponent(s) preferences. Accordingly, for a DM with a Nash profile, from some starting state s, two behaviours can occur: DM i stays at s or DM i moves away from s. In the first case, this means that there are no UIs for DM i from s, hence s is Nash stable for i. In other words, at a Nash stable state s, the states reachable from s are less than or equally preferred to s. Preference rankings in which $s \prec_i t$ for $t \in R_i(s)$ are thus removed from the possible rankings.

In the second case, this means that state t represents a UI from s, i.e., $s \prec_i t$. Preference rankings in which $s \succsim_i t$ can thus be removed from the possible rankings. Furthermore, this also means that t is greater than or equally preferred to the other states in the set of reachable states from s (i.e., t is the most or equally preferred UI). Preference rankings in which $t \prec_i u$ for $u \in R_i(s)$ can also be removed from the possible rankings. Algorithm 3 details the procedure for checking for sanctioning when the opponent has a Nash behavioural profile.

Since the DM is assumed to be a Nash player, they will stay at a Nash stable state and move away from states when there are UIs available. Note that there is no guarantee that the state moved to is Nash stable; it is simply a UI from the starting state. In other words, a DM moving from $startState$ to $endState$ is signalling that $endState$ is strictly more preferred to $startState$. Furthermore, the DM is also signalling that $endState$ is the most preferred state in the set of the DM's UMs from $startState$.

To summarize, the Nash profile behaves as follows:

- If DM i stays at s, remove rankings in which $s \prec_i t$ for $t \in R_i(s)$ because $t \precsim_i s$
- If DM i moves away from s to state t, remove rankings in which $s \succsim_i t$ because $s \prec_i t$ and those in which $t \prec_i u$ for $u \in R_i(s)$ because $t \succsim_i u$

Algorithm 3 isMoveSanctioned (Nash)($startState$, $endState$, $ranking$)

```
if endState ≠ startState then
    if endState ≾ startState then
        return True
    end if
end if
for state in unilateralMoves(startState) do
    if endState ≺ state then
        return True
    end if
end for
return False
```

4.4.3 GMR Profile

A state s is GMR stable for DM i iff for every $s_1 \in R_i^+(s)$ there exists at least one $s_2 \in R_{N-i}(s_1)$ with $s_2 \precsim_i s$ (Howard 1971). In other words, every UI that DM i has from s is sanctioned by an opponent UM; DM i anticipates that taking any one of its available UIs from s will result in a less preferred state once the opponent(s) counter-moves. Given this, and assuming that DMs are acting according to the behavioural profile that characterises this stability concept, there are two cases to consider: DM i stays at s or DM i moves away from s.

If DM i stays at s, this means that s is GMR or Nash stable for i, i.e., that either all of i's UIs from s are sanctioned by opponent UMs or there are no UIs. Thus, rankings which do not result in s being Nash stable (those in which $t \prec_i s$) or GMR stable state (those in which the UIs from s are not sanctioned by opponent UMs) can thus be removed from the set of possible rankings.

If DM i moves from s to t, this means that s is not GMR or Nash stable for i, i.e., there is at least one unsanctioned UI, t, from s. First, note that t is a UI from s, i.e., $s \prec_i t$. Any rankings in which $s \succsim_i t$ can thus be removed from the list of possible rankings. Next, one knows that i's movement to t is not sanctioned by opponent UMs; any rankings in which such a move is sanctioned can thus also be removed from the list of possible rankings. The procedure for verifying sanctioning for an opponent with a GMR profile is outlined in Algorithm 4.

The algorithm begins by checking whether $startState$ and $endState$ are the same. If this is the case (i.e., the DM decided to stay), the state is checked for GMR stability. If the start and end states are different, DM has signalled that $endState$ is a UI from $startState$. In other words, the move from $startState$ to $endState$ is not sanctioned. To check sanctioning for a GMR player in this case, it is necessary to verify that opponent moves from the UIs are less than or equally preferred to the starting state.

To summarize, the GMR profile behaves as follows:

- If DM i stays at s, remove rankings in which states in $R_i^+(s)$ are not sanctioned by opponent UMs
- If DM i moves away from s to t, remove rankings in which $t \precsim_i s$ because t should be a UI from s and those in which opponents can move from t to a state which is less preferred

Algorithm 4 isMoveSanctioned (GMR)($startState, endState, ranking$)

```
if startState = endState then
    for state in UIs(startState) do
        if not isMoveSanctioned (GMR)(startState, state, ranking) then
            return True
        end if
    end for
    return False
end if
if endState ≲ startState then
    return True
end if
for state in opponentsUMs(endState) do
    if state ≲ startState then
        return True
    end if
end for
return False
```

to s

4.4.4 SEQ Profile

A state s is SEQ stable for DM i iff for every $s_1 \in R_i^+(s)$ there exists at least one $s_2 \in R_{N-i}^+(s_1)$ with $s_2 \precsim_i s$ (Howard 1971). This definition is identical to GMR stability with the exception that the set of opponent UIs, rather than opponent UMs is considered. This case is much more difficult to implement as an inverse engineering algorithm since one assumes that an SEQ DM is aware of each of the other DMs' preferences.

Thus, state s is SEQ stable for DM i if every UI that DM i has from s is credibly sanctioned by an opponent. Once gain, DM i can stay at s or move away from s.

If DM i stays at s, this means that s is SEQ or Nash stable for i, i.e., that either all of i's UIs from s are sanctioned by opponent UIs or there are no UIs available. Rankings which do not result in s being Nash stable (those in which $t \prec_i s$) or SEQ stable (those in which the UIs from s are not sanctioned by opponent UIs) are removed from the set of possible rankings.

If DM i moves from s to t, this means that s is not SEQ or Nash stable for i, i.e., there is at least one unsanctioned UI, t, from s. First, note that t is a UI from s, i.e., $s \prec_i t$. Any rankings in which $s \succsim_i t$ can thus be removed from the list of possible rankings. Next, one knows that i's movement to t is not sanctioned by opponent UIs; any rankings in which such a move is credibly sanctioned can also be removed. The procedure for verifying sanctioning for an opponent with an SEQ profile is quite similar to that for the GMR profile and is outlined in Algorithm 5.

The algorithm first checks whether $startState$ and $endState$ are the same. If this is the case (i.e., the DM decided to stay), the state is checked for SEQ stability. If the start and end states are different, this means that $endState$ is a UI from $startState$. To check sanctioning for an

Algorithm 5 isMoveSanctioned (SEQ)($startState, endState, ranking$)

```
if startState = endState then
    for state in UIs(startState) do
        if not isMoveSanctioned (GMR)(startState, state, ranking) then
            return True
        end if
    end for
    return False
end if
if endState ≾ startState then
    return True
end if
for state in opponentsUIs(endState) do
    if state ≾ startState then
        return True
    end if
end for
return False
```

SEQ player in this case, it is necessary to verify that opponent UIs from the focal DM's UIs are less than or equally preferred to the starting state.

To summarize, the SEQ profile behaves as follows:

- If DM i stays at s, remove rankings in which states in $R_i^+(s)$ are not sanctioned by opponent UIs
- If DM i moves away from s to t, remove rankings in which $t \precsim_i s$ because t should be a UI from s and those in which opponents can move from t to a state which is less preferred to s

It should be noted that the SEQ profile is applicable only to conflicts with two DMs. The reason for this is that the set opponentsUIs($endState$) is clearly defined only when the UIs are known for all of the opponents. Although the observing DM's preference ranking is known to all of the other DMs, there is no guarantee that they are aware of each others' preferences. This makes the set opponentsUIs($endState$) difficult to ascertain. The DM with an SEQ profile checks whether its UIs are sanctioned by UIs from both the observing DM and from any other DM in the conflict; in order to know whether the sanctioning move is a UI, knowledge of all preferences is required.

4.4.5 SMR Profile

A state s is SMR stable for DM i iff for every $s_1 \in R_i^+(s)$ there exists at least one $s_2 \in R_{N-i}(s_1)$ such that $s_2 \precsim_i s$ and $s_3 \precsim_i s$ for every $s_3 \in R_i(s_2)$ (Howard 1971). The DM will stay at state s if, even after being able to respond to opponent sanctions, the resulting state(s) is/are less preferred to s. As usual, two things can occur: DM i remains at s or DM i moves away from s. In the former case, s is SMR or Nash stable. As was the case with GMR stability, one cannot know why i is choosing to stay; it could be that i has no UIs from s or that i's move to a UI is SMR sanctioned. Rankings which are compatible with this hypothesis are kept; those which are incompatible are discarded.

If DM i moves, then s is not SMR stable for DM i, i.e., there is at least one unsanctioned UI, t, from s. First, note that t is a UI from s; rankings in which $s \lesssim t$ can be removed from the list of possible rankings. Next, note that i's move to t is unsanctioned; any rankings in which this move is sanctioned can therefore also be removed from the list of possible rankings. Finally, the DM's should have a counter-move available which is more preferred to the starting state.

Algorithm 6 isMoveSanctioned (SMR)($startState$, $endState$, $ranking$)

```
if startState = endState then
    for state in UIs(startState) do
        if not isMoveSanctioned (SMR)(startState, state, ranking) then
            return True
        end if
    end for
    return False
end if
if endState ≲ startState then
    return True
end if
for state in opponentUMs(endState) do
    counterMoveExists = False
    if state ≲ startState then
        for counterMove in counterMoves(state) do
            if counterMove > startState then
                counterMoveExists = True
            end if
            if counterMoveExists = False then
                return True
            end if
        end for
    end if
end for
return False
```

The algorithm begins by checking whether $startState$ and $endState$ are the same; if this is the case, the state is checked for SMR stability.

If $startState$ and $endState$ are different, it is necessary to check whether the move from $startState$ to $endState$ is SMR sanctioned. The move is SMR sanctioned if there exists a state t in the opponents' UMs from $endState$ satisfying (1) $t \lesssim startState$ and (2) every counter move from t is less than or equally preferred to $startState$.

To summarize, the SMR profile behaves as follows:

- If DM i stays at s, remove rankings in which states in $R_i^+(s)$ are not sanctioned by opponent UMs once DM i has the chance to counter-move
- If DM i moves away from s to t, remove rankings in which $t \precsim_i s$ and those in which the move from s to t is sanctioned after the opponents move to a UM and DM i has the chance to counter-move

4.5 Advice Function

Once observations have been logged and preference rankings updated, this information is highly useful to the focal or observing DM, who can use them to make more informed decisions. The ADVICE function is called by the observing DM who would like information on how to proceed from the current state s of the conflict. Upon calling on the ADVICE function, the observing DM receives three pieces of information for s and for each of their UMs from s:

- **Expected payoff value of the state:** gives an (ordinal) expected payoff for s and for each of the observing DM's reachable states from s. The expected value takes a weighted average of the focal DM's payoff values based on all of the possible subsequent opponent moves. This value gives the DM a measure of the expected satisfaction for the state once all possible opponent responses have been considered.
- **Counter-move probabilities:** for each state, gives a probabilistic appraisal of how the opponents might react. In other words, given s or one of the observing DM's reachable states, the observing DM can then see how their opponents are likely to counter-move from that state. This tells the focal DM how the conflict is likely to progress once their move from state s has been made.
- **Opponent reachable states:** lists the states to which the opponents can move from each of the observed DM's available states, including the option to remain. In other words, for s and for the reachable states from s, the opponent's set of reachable states is included. This reminds the observing DM of which states the observed DM is theoretically capable of reaching.

4.5.1 Algorithms

The advice function given by Algorithm 7 provides the focal DM with several important pieces of information. Given the current state, ADVICE checks the focal DM's unilateral moves from that state. For each unilateral move, the algorithm loops through the opponents and through each possible preference ranking for the opponent and calculates the opponent's best response to the unilateral move. Once the function BESTRESPONSES is run, the sample size is incremented by the number of best responses in the output (there may be more than one best response). The array $numOccurrence$ records how many times each feasible state appears as a best response. The probability of each feasible state occurring at the immediate next step is calculated by dividing the number of occurrences of each feasible state by the total sample size.

Algorithm 7 ADVICE($currentState$)

```
for state ∈ unilateralMoves do
    numOccurrences = [0,0,...,0]                    ▷ One zero in the list for each feasible state
    sampleSize = 0
    for opponent ∈ opponentList do
        for rankings ∈ possibleRankings do
            responses = BESTRESPONSES(state, ranking)
            sampleSize ← sampleSize + length(responses)
            for numOccurrence ∈ numOccurrences do
                Add number of times the feasible state occurs in responses to numOccurrence
            end for
        end for
        probabilities = numOccurrences/sampleSize       ▷ List obtained by dividing each numOccurrence by the sample size
        expectedValue = expected value of payoff using probabilities
        Add state, its expected value and its probability to adviceStates
    end for
    return adviceStates
end for
```

The focal DM's expected value for each state is then calculated based on the probabilities and on the DM's own preference ranking. For example, if the probabilities for states s_1, s_2, s_3 and s_4 are 30%, 20%, 10% and 40%, respectively and the focal DM's preference ranking is $s_4 \succ s_1 \succ s_2 \sim s_3$ with associated payoffs 4, 3, 2, and 2, then the expected values are $E(s_1) = 3 \cdot 0.3 = 0.9$, $E(s_2) = 2 \cdot 0.2 = 0.4$, $E(s_3) = 2 \cdot 0.1 = 0.2$, and $E(s_4) = 4 \cdot 0.4 = 1.6$

The key step in the above algorithm is determining the opponent's best responses. These are moves that the opponent makes which are consistent with that opponent's behavioural profile. As such, opponents with divergent behavioural profiles will have different best response sets. For example, an opponent with a Nash behavioural profile is not concerned about possible sanctioning, while a GMR opponent will not move to states that could be sanctioned; the Nash DM's set of best responses thus includes states which the GMR's set does not. Algorithm 8 shows how the best response is computed. The heart of this algorithm is calling the function ISSANCTIONED, whose Boolean output (True/False) depends on the behavioural profile.

Algorithm 8 BESTRESPONSES($state$, $ranking$)

```
bestResponses ← [ ]
for state in unilateralMoves do
    if not ISSANCTIONED then
        add outcome to bestResponses
    end if
end for
return bestResponses
```

The opponent's behavioural profile is therefore useful at two steps in the overall process. First, it assists with the removal of infeasible preferences; second, it allows for a more exact determination of best responses, which eventually contributes to providing useful advice. The use of behavioural profiles in the advice stage ensures that expected values and probabilities truly reflect how the conflict could unfold; including sanctioned states in an opponent's best response list would skew the results.

4.5.2 Complexity

Given a conflict with s feasible states, the number of possible rankings of those states is given by the ordered Bell number (Good 1975) which is approximated by $Bell(s) \approx \frac{s!}{2(\ln 2)^{s+1}}$ (Gross 1975). Ordered Bell numbers quickly become quite large and, as such, are the source of limitations regarding the space and time complexity of the algorithms presented thus far.

With respect to storage space, saving the list of all possible rankings is simply not feasible for conflicts with 12 states or more. Assuming that four bits are used to store the numerical ranking of a state, up to 16 states can be ranked; if a conflict has more than 16 feasible states, 8 bits are needed to store each state's ranking. Table 4.1 shows the space required to store the number of possible rankings for conflicts with 8 - 17 feasible states. As the table illustrates, the storage space needed quickly reaches the order of terabytes for conflicts with 13 states, petabytes for conflicts with 15 states, and exabytes for conflicts with 17 states. The amount of storage needed could be reduced by further compressing the data; however, reading the data from the compressed file is slower than computing the possible rankings from scratch. Thus, rather than attempting to store any of the data, the possible rankings should be generated from scratch at each iteration, then discarded.

Number of feasible states	Number of possible rankings	Storage Requirement
8	545 835	2 MB
9	7 087 261	32 MB
10	102 247 563	511 MB
11	1 622 632 57 3	9 GB
12	28 091 567 595	169 GB
13	526 858 348 381	3 TB
14	10 641 342 970 443	74 TB
15	230 283 190 977 853	2 PB
16	5 315 654 681 981 355	43 PB
17	130 370 767 029 135 901	2 EB

Table 4.1: Storage space requirements for number of feasible states, assuming 4 bits are used to store each state's ranking value. The storage requirement quickly increases with the number of feasible states.

With the number of opponents as constants and the number of states in a DM's set of UMs limited to at most s, three additional factors affect time complexity. First and foremost, $O\left(\frac{s!}{\ln 2^s}\right)$ time is required to loop through the list of possible rankings. Second, it is necessary to cycle through the focal DM's UMs, which adds at most a factor of s. Third, time is required to verify whether a ranking is valid or not given the opponent's behavioural profile. This time depends on the opponent's behavioural profile: $O(1)$ time is required for the Default profile; $O(s)$ time is required for the Nash profile; $O(s^2)$ time is required in the worst case for the GMR profile; and $O(s^3)$ is required in the worst case for the SMR profile. Thus, in the worst case, the algorithm runs in $O\left(\frac{s^4 \cdot s!}{\ln 2^s}\right)$ time.

As a reference point, running the ADVICE function on two observed moves for a conflict with 11 states on a typical home computer took between 19 hours and 10 minutes and 59 hours and 40 minutes depending on the opponent profile. Although many conflicts studied using the Graph Model have 11 feasible states or less, many others have 12 states or more.

Finding ways to optimise the performance of the ADVICE function remains a key research step. Some ideas under consideration include implementing parallel computing, using specialty research servers, and removing certain infeasible rankings from the outset. For example, it may be known that a certain state is (or is not) a given DM's most/least preferred state; in this case, any rankings which do not satisfy this condition need not be considered.

4.5.3 Special Cases

Before moving on to the first case study, it will be useful to highlight cases in which the inverse engineering algorithms might return an empty set of feasible rankings for one or more opponents. One reason for this is that the behavioural profile assigned to the opponent is incorrect; as such, the opponent's behaviour is inconsistent with the behavioural profile, which leads to the eventual removal of all feasible rankings. This situation occurs in the case study presented in Section 4.6.

Next, an empty set of feasible states could be indicative of a change in preferences. Thus, the behaviour observed at a given point in the conflict could potentially be at odds with the DM's new preference ranking and with their subsequent actions.

Finally, opponents who purposely deviate from their behavioural profiles (e.g. in an attempt at trickery) could render the set of feasible rankings empty. Once again, this is because their observed behaviour would, at times, be consistent with their true preferences rankings, and at times be inconsistent with it.

As the discussion above has shown, there is something of a "safeguard" built into the inverse engineering methodology. Although it is essentially a true/false flag, knowing that the set of feasible preference rankings of an opponent is empty helps the observing DM realise that something is amiss. Diagnosing the problem is not necessarily straightforward, but the DM at least is cognisant of the fact that additional verification is required: it may be that the opponent's preferences have changed over time, that the behavioural profile assigned to them is incorrect, or that the opponent has attempted to engage in deception.

4.6 Application

Consider the 2-DM conflict model described in (Hipel 2001). Developers (D) heading a project are monitored by environmentalists (E). The developers can choose sustainable development or not. Environmentalists, on the other hand, can be proactive or reactive.

Table 4.2: DMs and outcomes for sustainable development conflict

Environmentalists				
Proactive	1	1	0	0
Developer				
Sustainable Development	1	0	1	0
State	s_1	s_2	s_3	s_4

The Graph Model representation of this conflict is shown in Figure 4.1 where the solid lines represent the environmentalists' UMs and the dashed lines correspond to the developer's UMs. DM moves are assumed to be reversible and thus are represented by bidirectional arcs. Notice, for example, that the developers have a UM from state s_1 to state s_2 by changing its option selection from practising sustainable development (a 1 is opposite this option in Table 4.2 for state s_1) to not choosing sustainable development (as indicated by a 0 opposite sustainable development for s_2), while the environmentalists do not change their option choice (a 1 is placed opposite the option proactive for both states s_1 and s_2). This UM for the developers is portrayed by the dashed arrow line going from state s_1 to state s_2 at the top of Figure 4.1.

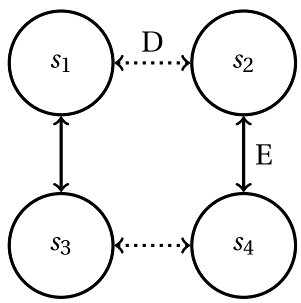

Figure 4.1: Graph Model of the sustainable development conflict

4.6.1 Starting Point

At the beginning of the conflict, the environmentalists have no information about the developer's preferences and are thus unable to exclude any preference rankings outright. Given the number of options, there are $Bell(4) = 75$ possible preference rankings for the developers (from Equation 4.1). It is assumed that the environmentalists' preferences are known to be $s_1 \sim_E s_2 \succ_E s_3 \succ_E s_4$ (with corresponding payoffs 4, 4, 3, and 2, respectively). The analysis is being conducted from the environmentalists' point of view.

The environmentalists are interested in discovering the developer's true preference ranking, which is $s_2 \succ_D s_1 \succ_D s_3 \succ_D s_4$. The conflict begins at status quo state s_4 in which neither DM has selected their option. Suppose that the developer moves first and takes the conflict into state s_3 by selecting the option for sustainable development. Upon observing this move, the environmentalists apply the inverse engineering and advice algorithms (1st iteration). Next, the environmentalists respond by moving from state s_3 to state s_1. Then, the developers move the conflict from state s_1 to state s_2; this move is once again observed by the environmentalists (2nd iteration). Once the conflict at in state s_2, neither DM makes any moves (3rd iteration).

4.6.2 General Observations

For reference, Table 4.3 summarizes how many rankings remain once each algorithm has removed infeasible preference rankings.

Table 4.3: Number of rankings remaining after iterations of the algorithms for developer ranking

Profile	1st iteration	2nd iteration	3rd iteration
Default	44	26	26
Nash	31	13	13
GMR	19	0	0
SMR	25	9	9

First, some general observations: at each of the first two iterations, regardless of behaviour profile, the set of possible preference rankings is reduced; the final iteration yields no additional information on the possible preference ranking. This is not surprising given that by the third iteration, it is difficult for this small conflict to provide any new information.

When the developer's profile is set to 'default', less rankings are removed at each iteration compared to the other behaviour profiles. Since the 'default' profile is by design the most conservative in terms of ranking removal, this is to be expected. Given how the removal is executed in this case, rankings in which $s_3 \prec s_4$ and $s_2 \prec s_1$ are the only ones to be removed.

For the Nash profile, the first iteration removes more than half of the possible rankings; not only are rankings in which $s_3 \prec s_4$ excluded, but also those in which $s_3 \sim s_4$ since a Nash DM being indifferent between two states will not move from one to the other. At the end of the second iteration, only 13 possible rankings remain.

If the developer is assumed to have a GMR behavioural profile, the first iteration of the algorithms removes more rankings than any other profile. This is because the developers have provided quite a bit of information: s_4 is not a GMR stable state, which means that s_3 is a UI from s_4; furthermore, s_4 is an unsanctioned UI. After the second iteration, however, there are no longer any possible preference rankings remaining. This means that the developer's behaviour is not consistent with a GMR profile; in other words, the assumption that the developers make decisions based on GMR stability is incorrect.

To see why this is the case, note that the developers' move from s_1 to s_2 means that s_2 is an unsanctioned UI from s_1. For the GMR profile, it is necessary to check to which state the environmentalists can bring the conflict from s_2 by considering their UMs. From s_2, the environmentalists have a UM to s_4; s_4 is the developers' least preferred state, hence their UI from s_1 to s_2 is GMR sanctioned by the threat of s_4. Thus, assuming that the developer is a GMR player, the move from s_1 to s_2 should not take place; the fact that it does means that the developers are not behaving consistently with a GMR profile. Thus, this information is valuable

Table 4.4: Remaining preference rankings after three iterations

Default Profile	**Nash Profile**	**SMR Profile**	**GMR Profile**
$s_3 \succ s_2 \succ s_1 \succ s_4$	✓	✓	
$*\mathbf{s_2 \succ s_1 \succ s_3 \succ s_4}$	✓		
$s_2 \succ s_3 \succ s_1 \succ s_4$	✓	✓	
$s_3 \succ s_2 \succ s_4 \succ s_1$	✓	✓	
$s_2 \succ s_3 \succ s_4 \succ s_1$	✓	✓	
$s_3 \succ s_4 \succ s_2 \succ s_1$	✓		
$s_3 \succ s_2 \succ s_4 \sim s_1$	✓	✓	
$s_2 \succ s_1 \succ s_3 \sim s_4$			
$s_2 \succ s_3 \succ s_4 \sim s_1$	✓	✓	
$s_3 \succ s_4 \succ s_1 \sim s_2$			
$s_3 \succ s_2 \succ s_1 \sim s_4$			
$s_2 \succ s_3 \sim s_1 \succ s_4$	✓		
$s_3 \succ s_4 \sim s_2 \succ s_1$	✓		
$s_2 \succ s_4 \sim s_3 \succ s_1$			
$s_3 \succ s_4 \sim s_2 \sim s_1$			
$s_2 \succ s_4 \sim s_3 \sim s_1$			
$s_2 \sim s_1 \succ s_3 \succ s_4$			
$s_2 \sim s_3 \succ s_1 \succ s_4$	✓	✓	
$s_2 \sim s_3 \succ s_4 \succ s_1$	✓	✓	
$s_4 \sim s_3 \succ s_2 \succ s_1$			
$s_2 \sim s_1 \succ s_4 \sim s_3$			
$s_2 \sim s_3 \succ s_4 \sim s_1$	✓	✓	
$s_4 \sim s_3 \succ s_2 \sim s_1$			
$s_2 \sim s_3 \sim s_1 \succ s_4$			
$s_4 \sim s_3 \sim s_2 \succ s_1$			
$s_4 \sim s_3 \sim s_2 \sim s_1$			

* Developer's true preference ranking

insofar as it provides insight into the developer's behaviour: one can rule out the possibility that the developers conform to the GMR profile.

Finally, assuming that the developers have an SMR profile, 25 rankings are removed at the first iteration. The second iteration, however, leaves only nine possible preference rankings. The developer's movement from a given state signals that the state from which it moved is not SMR; a UI is available either because it is unsanctioned by opponent UMs or since the focal DM is able to escape the sanction.

Table 4.4 shows the rankings that remain after the three iterations for each behaviour profile. For simplicity, the 26 rankings remaining in the default profile are listed in full; those which remain for Nash, GMR, or SMR are subsets of the list of 26 and are denoted by a checkmark (✓) in the appropriate column. The developer's true preference ranking is highlighted in bold and preceded by an asterisk (*).

Note that the developer's true preference ranking is contained in the lists generated by both the default and Nash profile algorithms. This indicates that, given its true preferences, the developer's behaviour is consistent with both a default and Nash behavioural profile. This is not, however, the case with SMR: the developer's true preferences are not present in the list generated by the algorithm. If it is assumed that the developers always move in accordance to their behaviour profiles, this absence means that the developers are not SMR DMs. In fact, it is simple to verify that if the developers were truly SMR DMs, they would not have moved from s_1 to s_2: such a move is sanctioned by an opponent move to state s_4, from which the only possible counter-move to s_3 is less preferred to s_1.

Thus, in this case, the environmentalists have mistakenly attributed an SMR profile to the developers and although the algorithm generates a non-empty list of possible rankings, the true ranking is not among them. This emphasizes the importance of attributing a proper behavioural profile to the DM whose preferences one wishes to uncover. While some profiles may eventually be discarded once they generate an empty list of rankings, this may not always occur; it remains the responsibility of the modeller to carefully determine the appropriate behavioural profile.

4.6.3 Advice Function

Given a list of possible opponent rankings, the next step is to use this information to improve one's own decision-making: the ADVICE function is created in order to facilitate this process. Just like their opponent(s), the focal DM also has its own behavioural profile. The advice function provides counsel based on a range of behavioural profiles, leaving the final decision to the DM.

Table 4.5 summarizes the results obtained for each of the behavioural profiles that could be assigned to the developers. Based on these results, the environmentalists should move to state s_1 regardless of the developers' behavioural profile if they wish to maximise their payoff for the immediate next step. In all cases except SMR, the developer is most likely to remain at state s_1 in response.

When the developers are assumed to have a default behavioural profile, their move from state s_4 to state s_3 reduces the number of possible rankings from an initial 75 to 44. From state s_3, the environmentalists can either remain (top sub-row of Default row of Table 4.6) or move to state s_1 (bottom sub-row of Default row of Table 4.6). In the former case, the expected value is 2.77. Based on the environmentalists' preference ranking which assigns a payoff of 4 to states s_1 and s_2, a payoff of 3 to state s_3, and a payoff of 2 to state s_4, remaining at s_3 is expected to fall slightly below the payoff value of state s_3. This is because at the next immediate step, the developers have a 22.8% chance of moving from s_3 back to s_4, but a 77.2% chance of staying at s_3.

If, however, the environmentalists were to move from state s_3 to state s_1 rather than stay at s_3, the expected value is now 4; based on the possible developer preference rankings, there is a 37.1% chance that the developers will move from s_1 to s_2 and a 62.9% chance that the developers will stay. Either case is acceptable to the environmentalists as states s_1 and s_2 are tied as their most preferred state. The last column of the table lists which states are reachable for the opponent (in this case, the developers) from states s_3 and s_1, respectively.

If the developers are assumed to have a Nash profile, 31 possible rankings remain once the environmentalists update the list. As before, the environmentalists can choose to either

Table 4.5: Environmentalists' ADVICE function from s_3 after observing the developers move from state s_4 to state s_3.

Profile	Number of possible rankings (1st iteration)	Expected payoff values	Counter-move probabilities	Developer's reachable states
Default	44	s_3: 2.77	s_4 = 22.8% s_3 = 77.2%	$\{s_3, s_4\}$
		s_1: 4	s_2 = 37.1% s_1 = 62.9%	$\{s_1, s_2\}$
Nash	31	s_3: 3	s_3 = 100%	$\{s_3, s_4\}$
		s_1: 4	s_2 = 41.9% s_1 = 58.1%	$\{s_1, s_2\}$
GMR	19	s_3: 3	s_3 = 100%	$\{s_3, s_4\}$
		s_1: 4	s_1 = 100%	$\{s_1, s_2\}$
SMR	25	s_3: 3	s_3 = 100%	$\{s_3, s_4\}$
		s_1: 4	s_1 = 36% s_2 = 64%	$\{s_1, s_2\}$

remain at state s_3 or move to state s_1.

In the first case, they can be assured that the conflict will remain at state s_3. Indeed, by examining the developer's preference ranking, one can see why this is the case: state s_3 is in fact Nash stable for the developers since there are no available UIs. Thus, the environmentalists can see that although the developers can, in theory, move the conflict to state s_4, this is unlikely to occur. This type of information is useful to the focal DM who can see whether some opponent moves are indeed possible. Knowing that certain opponent moves have a zero or low probability of occurring could help overcome sanctioning concerns. The environmentalists can see that staying at state s_3 is "safe": the developers will not move back to state s_4, which is the environmentalists' least preferred state.

In the second case, the environmentalists can move from state s_3 to state s_1. As in the situation with the unknown profile, the environmentalists are guaranteed a payoff of 4 since they are indifferent between the developers moves to s_1 and to s_2.

If the developers are assumed to have a GMR behavioural profile, observing their initial move allows the list of possible rankings to be reduced to 19. Based on the remaining rankings, the expected value of the environmentalists' decision to stay at s_3 is 3, while that of moving to state s_1 is 4. In this situation, the developer's next move is known with certainty in both cases: the developer will stay at s_3 and stay at s_1 even though there are other moves available.

Finally, if the developers are assumed to have an SMR behavioural profile, the number of rankings is reduced to 25. If the environmentalists decide to remain at state s_3, their expected value is 3, and the developers will remain at state s_3. If the environmentalists move from state s_3 to state s_1, they can expect a payoff of 4.

Based on the information provided by the ADVICE function, the environmentalists would make the same decision regardless of the developer's behavioural profile. A move to s_1 would, in all cases, yield a payoff of 4 at the next immediate play. Now assume that the environmentalists take this advice and move from state s_3 to state s_1 and that the developers respond by moving from state s_1 to state s_2. Once the developer's possible rankings have been updated, the environmentalists once again seek advice.

The results of the next call of the ADVICE function are shown in Table 4.6. Based on these results, the environmentalists should stay at s_2 if they wish to maximise their payoff for the immediate next step, regardless of the developers' behavioural profile. If the environmentalists remain at s_2, the developers will certainly remain if they are either Nash or SMR and are very likely to remain if they are Default. The results also show that the sequence of moves by the developer is inconsistent with a GMR behavioural profile.

Table 4.6: Environmentalists' ADVICE function from state s_2 after observing the developers move from state s_1 to state s_2.

Profile	**Number of possible rankings (2nd iteration)**	**Expected payoff values**	**Counter-move probabilities**	**Developer's reachable states**
Default	26	s_4: 2.5	s_4 = 50% s_3 = 50%	$\{s_3, s_4\}$
		s_2: 4	s_2 = 76.5% s_1 = 23.5%	$\{s_1, s_2\}$
Nash	13	s_4: 3	s_3 = 100%	$\{s_3, s_4\}$
		s_2: 4	s_2 = 100%	$\{s_1, s_2\}$
GMR	0			
SMR	9	s_4: 3	s_3 = 100%	$\{s_3, s_4\}$
		s_2: 4	s_2 = 100%	$\{s_1, s_2\}$

In all cases, the number of possible rankings has decreased. Additionally, the number of possible rankings that remain depends on the behavioural profile assigned to the developers; a move may provide more or less information depending on the profile. Note that the number of possible states for the GMR behavioural profile is zero; this reveals that the developer's actions are inconsistent with the GMR profile.

When the developer is assumed to have a default profile, the environmentalists can expect a higher payoff by remaining at state s_2 rather than moving to state s_4. This is because staying at state s_2 leaves the developers able to move only to state s_1 which is tied as the environmentalists' most preferred state. If the developers are assumed to have a Nash or SMR profile, the expected payoff is higher if the environmentalists remain at s_2. In both cases, the developer's next moves are deterministic; the environmentalists can be sure of where the conflict will go once they choose to move.

4.7 Chapter Summary

An inverse approach to preference elicitation in the Graph Model can be a useful tool for DMs and analysts. Not only does it provide a way to more accurately gauge opponent preferences, it also assists in real-time decision-making. Specifically, the information furnished by the Advice function allows focal DMs to make improved strategic decisions based on up-to-date data. This chapter highlights the ideas and algorithms underpinning the Advice function, which relies on a DM's behavioural profile to determine which moves would not occur and, as a consequence, which states are likely to occur at the immediate next step. The expected value gives the focal DM an idea of the ordinal value of each of their possible moves, while the next state probabilities show the likelihood of the opponent's immediate next move.

As demonstrated by the sustainable development application, the Advice function is a useful decision-making tool; the focal DM is able to better assess the risks of moving away from a state versus staying at that state. A particularly helpful piece of information is whether the opponent is likely to, given its behavioural profile, actually move the conflict to all of its reachable states. If this is not the case, the focal DM gains additional insights about the possibility of sanctioning by the opponent.

Although this conflict is relatively small, the example clearly illustrates why the information provided by the Advice function is quite relevant. The environmentalists are able to choose their next move based on the expected value of that move and on how the conflict is likely to evolve at the immediate next step. The analysis provided aids focal DMs in making more enlightened decisions; depending on their levels of risk aversion, different focal DMs might choose different actions.

Chapter 5

Kinder Morgan Case Study

5.1 Introduction

As the country with the third largest oil reserves in the world and as the world's sixth largest producer, Canada depends on oil extraction and refining processes for its economic development (National Energy Board 2015a). Oil sands make up 90% of Canada's reserves, with conventional oil accounting for the remaining 10% (National Energy Board 2015b). The most important crude oil reserves in Canada are located in the Western Canada Sedimentary Basin, which traverses parts of Yukon, Northwest Territories, British Columbia, Alberta, and Saskatchewan. The oil is mainly produced in Alberta, which generated 77% of the country's production in 2013 (National Energy Board 2014). The Alberta oil sands alone have an estimated 1.8 trillion barrels of oil in place, of which an estimated 168 billion are ultimately recoverable (Natural Resources Canada 2013).

A net oil exporter, Canada has historically exported most of its oil to the United States. Recently, however, Canada has pushed for market expansion overseas to regions such as South America, Europe, and Asia (National Energy Board 2015a). To this end, the government of Canada has been trying to make use of the country's extensive pipeline system, which covers over 35,000 kilometres, to transport crude oil originating from Alberta's oil sands to the coasts for eventual shipping overseas.

The transportation of crude oil is primarily done using Canada's pipeline system which has lines transporting domestic crude oil to refineries and to the United States and transporting imported crude oil to refineries (Natural Resources Canada 2014). In 2014, pipelines transported more than seven times the crude oil exports than marine, rail, and trucks combined (National Energy Board 2015b). Due to their scope, which often crosses provincial or international boundaries, Canadian pipeline construction and expansion projects have been surrounded by controversy.

The Trans Mountain Expansion Project (TMEP), managed by Trans Mountain Pipeline Corporation, a wholly-owned subsidiary of the Texas-based Kinder Morgan Inc, is designed to expand the existing Trans Mountain pipeline that runs from Strathcona County, Alberta to marketing terminals and refineries in Burnaby, British Columbia on Canada's west coast (Trans Mountain Pipeline ULC 2013). This project is wide in scope as it crosses many provincial and regional boundaries within the country; consequently, it is currently at the centre of a conflict involving a range of stakeholders including municipalities, provincial governments, First Nations, Trans Mountain, and environmental groups.

The goal of this chapter is to analyse this dispute using an inverse engineering approach, which makes use of observable decision-maker (DM) actions to provide analysts with information regarding possible evolutions of the conflict. This information is valuable to analysts and DMs involved in the dispute who are deciding on their courses of action. As will be demonstrated, the data provided are not only useful, but are also well suited to interpretation by a variety of decision-making profiles.

5.2 Project Background

According to TMEP documents, the existing 1,150 km pipeline will be twinned, allowing for greater oil transporting capacity from the current 300,000 to 890,000 barrels per day (Trans Mountain Pipeline ULC 2013). Once completed, the new pipeline segment will carry heavy crude oils, while the existing pipeline will carry refined products, synthetic crude oils, and light crude oils. Approximately 980 km will be new pipeline and 193 km will be reactivated pipeline (Trans Mountain Pipeline ULC 2013). In addition to new pipeline, the project also involves the construction of 12 pumping facilities, 19 new storage tanks, and 3 new berths located in the Westridge Marine Terminal in Burnaby (Trans Mountain Pipeline ULC 2013; Kinder Morgan Canada Limited 2017b). For reference, a map of the TMEP is provided in Figure 5.1. The analysis of this conflict is based on events up until November 2017.

5.2.1 Canadian Regulatory Approval Process

According to the *Canadian Environmental Assessment Act, 2012 (CEAA)* and the *National Energy Board Act (NEBA)*, the National Energy Board (NEB) is responsible for issuing a certificate of public convenience and necessity for pipeline projects (Government of Canada 2014; *Canadian Environmental Assessment Act S.C. 2012, c.19, s.52* 2014). The certificate is issued based on whether the proposed project is in the public interest, which includes consideration of economic, social, and environmental impacts (Becklumb 2012). Section 52(4) of the *NEBA* sets a 15 month time limit from the receipt of a complete application to the NEB's recommendation report (Government of Canada 2014). Based on the NEB's assessments and recommendations, the Governor in Council "directs the NEB to issue a decision statement to the pipeline company informing it of the decision" (Becklumb 2012).

Typically, NEB board members are tasked with overseeing the review of each pipeline project in accordance with the Board's regulatory framework. This framework is composed of laws, requirements, and guidance: "[o]ur regulations, conditions and guidance start by defining the safety and security, environmental protection and economic efficiency outcomes (or performance objectives) to be achieved." (National Energy Board 2016c). Once a project has been approved, the NEB continues to be involved in the oversight and compliance verification of the project throughout its life cycle (National Energy Board 2016b).

In February 2018, the federal Liberal government announced new legislation to overhaul the environmental assessment process for major natural resources projects (*An Act to enact the Impact Assessment Act and the Canadian Energy Regulator Act, to amend the Navigation Protection Act and to make consequential amendments to other Acts (First Reading)* 2018). The

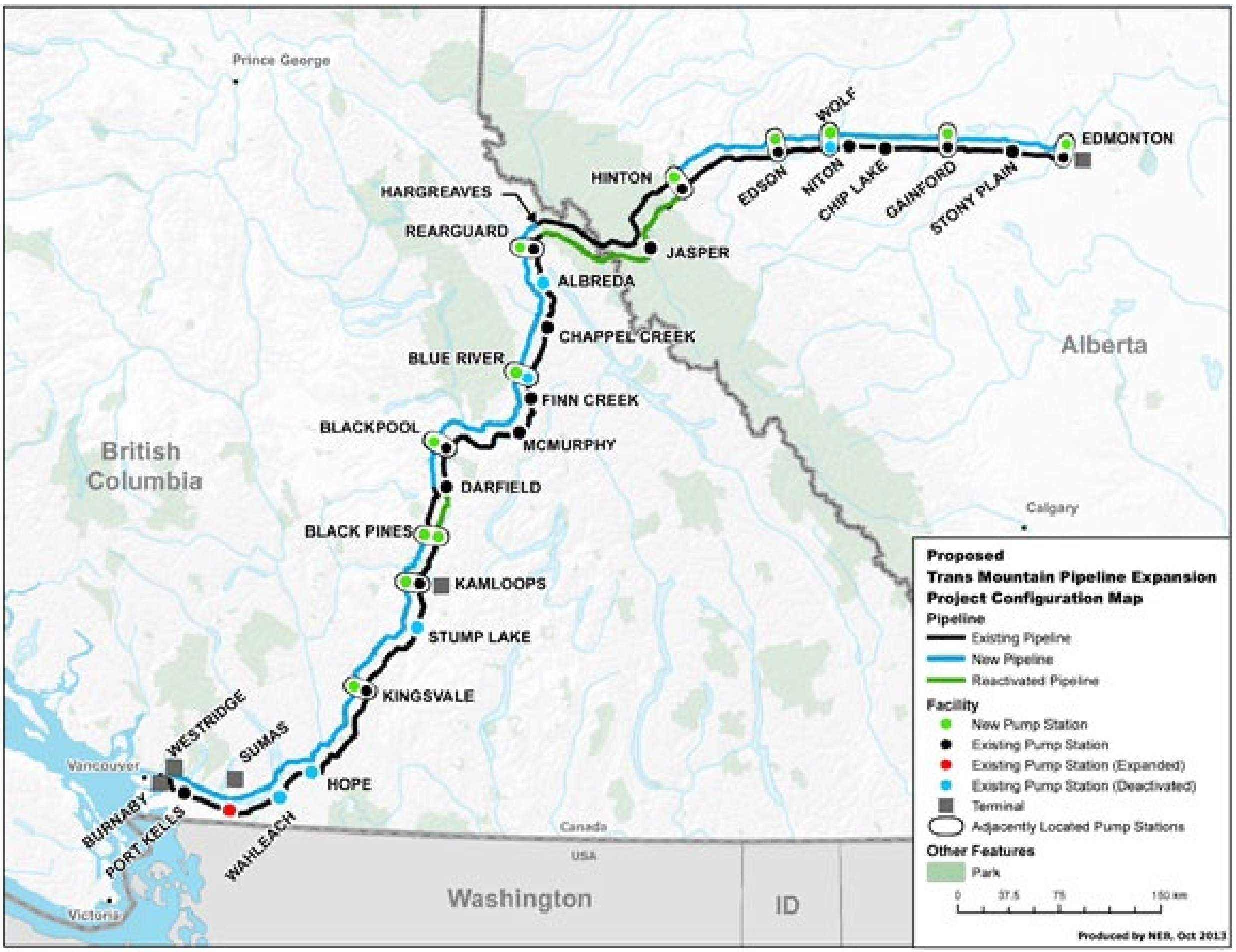

Figure 5.1: Trans Mountain Expansion Project Pipeline Route
Source: National Energy Board

bill, which has yet to pass into law, creates the Impact Assessment Agency of Canada, a single agency which would review large projects and evaluate their environmental, health, social, and economic impact. Pipeline projects in particular would be regulated by the Canadian Energy Regulator (CER) which replaces the NEB. When asked whether the TMEP would be approved under the new bill, Canada's Environment Minister Catherine McKenna responded in the affirmative (Tasker 2018).

On December 16, 2013, Trans Mountain Pipeline ULC filed its application with the NEB for the TMEP (Trans Mountain Pipeline ULC 2013). The application was considered complete on April 2, 2014 (National Energy Board 2016a); however the NEB's recommendation report was not issued until May 2016. The delay was due to Trans Mountain changing its preferred route corridor which resulted in additional hearings. Another source of delay was that one of the NEB board members dismissed some of Trans Mountain's filed evidence on oil market supply and demand which then required the collection of additional information from Trans Mountain and other intervenors (National Energy Board 2016a).

In its May 19, 2016 report, the NEB recommended the approval of the TMEP subject to 157 conditions regarding pipeline safety; emergency preparedness and response; environmental protection; ongoing consultation, particularly with Aboriginal communities; socio-economic matters; financial responsibility; and affirmation of commercial support (National Energy Board 2016a). On November 29, 2016 the Government in Council led by Prime Minister Justin Trudeau accepted the NEB's recommendations and directed the NEB to issue the Certificate of Public Convenience and Necessity for the TMEP (Governor in Council 2016; Office of the Prime Minister 2016).

Before Trans Mountain can begin full construction on the TMEP, it must file additional documentation to the NEB to demonstrate its compliance with the conditions imposed on the project (National Energy Board 2017a). Out of the 157 conditions, 98 necessitate additional documentation prior to the beginning of construction. Furthermore, the detailed pipeline route must go through the route approval process which is currently ongoing: many hearings were scheduled in the summer and fall of 2017, while other hearings have yet to take place (National Energy Board 2017c).

In addition to federal requirements, the TMEP is also subject to 37 conditions attached to the Environmental Certificate received from the Government of British Columbia. These conditions were issued on January 10, 2017 by the former Minister of Environment, Mary Polak, and the former Minister of Natural Gas Development, Rich Coleman (Polak and Coleman 2017). British Columbia's *Environmental Assessment Act* requires that the province release its own environmental assessment certificate for such projects, although it can rely on the NEB's process to arrive at its decision (*Environmental Assessment Act SBC 2002, c. 43* 2002). The conditions pertain to Aboriginal engagement, public engagement, construction, regulation, environment/socio-economics, and emergency response/preparedness (Government of British Columbia 2017), and are designed to supplement the conditions issued by the NEB (Polak and Coleman 2017).

5.2.2 Major Issues

This is a large-scale, inter-provincial pipeline project, and, like many others before it, brings several issues to the forefront. In recent years, pipeline projects have been the subject of many controversies and ideological disagreement between those touting the economic benefits and those concerned about the impacts on the environment and on First Nations. Furthermore, pipelines also illustrate differences in power between the three levels of government: federal, provincial, and municipal. Major issues related to the TMEP can be classified into four types: economic benefits, environmental concerns, duty to consult, and balance of power.

Economic Benefits

According to a report issued by the Conference Board of Canada, the TMEP would result in an estimated $18.5 billion in fiscal benefits during its development and over the first 20 years of operations (Burt 2016). It is also claimed that the project would generate 15,000 new jobs during construction as well as 440 permanent jobs per year during operation (Natural Resources Canada 2017). An analysis by the global consulting firm Muse Stancil commissioned by Trans Mountain and submitted to the NEB as part of the application process found that the total benefits of the project could reach $73.5 billion (in 2012 Canadian dollars) for Canadian crude oil producers (Earnest 2015).

This project is particularly important for the diversification of Canada's oil export markets, the underlying assumption being that diversifying Canada's oil market would result in higher prices for Canadian oil. The project would allow exportation to "Washington State and northeast Asia (Japan, China, South Korea and Taiwan) and to secondary markets in the United States such as California, Hawaii and Alaska" (Natural Resources Canada 2017). In its recommendation report, the NEB board found that "there would be a considerable benefit gained by providing Canadian shippers with more flexible and diverse markets, the ability to manage risk associated with competing in multiple markets, the ability to manage development and operational risk, and a likely reduction of discounts to Canadian crude" (National Energy Board 2016a).

These benefits have been challenged by several economists, notably by Robyn Allan and by David Hughes who argue that many of the stated benefits of the TMEP do not hold up to scrutiny. Both argue that the Canadian crude discount is not a discount, but a result of the transportation cost of Canadian oil (Allan 2016; Hughes 2017). According to Allan, there are no Asian markets for Alberta's oil (Allan 2016), while Hughes states that due to transportation costs, Canadian heavy oil would sell at a lower price than U.S. oil in Asia (Hughes 2017). Hughes also highlights the fact that Muse Stancil's report was based on several assumptions which no longer hold, including an overestimation of oil supply and that no other pipeline projects would be built (Hughes 2017).

Environmental Concerns

Environmental groups oppose the project based on concerns about the TMEP's impact on the environment and on climate change. The Raincoast Conservation Foundation and Living Oceans Society argues in its filing that the NEB's environmental assessment, which did not include an analysis of the impacts of marine shipping, violated the requirements of the *Species at Risk Act (SARA)*. Under the *SARA*, "If a wildlife species is listed as an extirpated species, an endangered species or a threatened species, the competent minister must prepare a strategy for its recovery" (*Species at Risk Act S.C. 2002, c. 29* 2014). Given that the TMEP would affect the habitat of killer whales along the coast, a submission of mitigation measures was required, but not submitted since the project's impact on marine shipping was not evaluated (Kung 2017).

The City of Burnaby and the City of Vancouver have been vocal opponents of the TMEP. In their court filings, both cities contend that the NEB failed the statutory requirements set out in the *CEAA* and in the *NEBA* (Kung 2017). More specifically, the City of Burnaby argues that the NEB did not consider alternate locations for one of the marine terminals and that the NEB did not collect evidence regarding the project risks until after the GIC approved the project. The City of Vancouver's claims are centred around the lack of proper assessment of marine activities, of the effects of greenhouse gas emissions, and of proper spills assessment.

Trans Mountain argues that transporting oil via pipelines rather than shipping it in tankers trucks or rail cars results in "safer, more efficient and more economic shipment of oil between Alberta and BC" (Trans Mountain 2017).

Duty to Consult and Land Rights

Under the Canadian Constitution, Aboriginal peoples' rights are protected; however, there are no clear guidelines surrounding consent (Nosek 2017). As such, Indigenous opposition or lack of consent on major projects such as pipelines "may or may not have an impact on whether the project proceeds" (Boreal Leadership Council 2015). In fact, the Crown's duty is limited to consultation with Indigenous communities and does not extend to reaching agreement with these communities on projects that could significantly impact them.

What's more, the Crown is legally allowed to delegate parts of the consultation to project proponents (i.e., to pipeline companies) and to "discharge its duty to consult and accommodate through the environmental assessment process" (Nosek 2017). This typically results in Indigenous communities having little control or input over natural resource development projects.

First Nations opposition to pipeline projects has taken on many different forms, including issuing declarations and filing lawsuits. Declarations such as the Coastal First Nations Declaration (*Coastal First Nations Declaration* 2010), the Save The Fraser Declaration (Indigenous Nations of the Fraser River Watershed 2010), the International Treaty to Protect the Salish Sea (*International Treaty to Protect the Salish Sea* 2014), and Treaty Alliance Against Tar Sands Expansion (*Treaty Alliance Against Tar Sands Expansion* 2015) have been signed by hundreds of First Nations.

In addition to issuing declarations asserting their land rights, First Nations are also employing legal tools to challenge the NEB's initial decision to approve the project. Broadly speaking, First Nations are opposed to the project on the grounds that the NEB's consultation was fundamentally inadequate and that, as a result, the decision to recommend the TMEP for approval was not truly made in the public interest. As Elin Sigurdson, counsel for Upper Nicola Band, summarized in the opening argument: "[a] decision by the government that is not procedurally fair, that contravenes legislative requirements, and that is contrary to the honour of the Crown and its constitutional duty to advance reconciliation with Indigenous peoples – is not in the public interest" (Kung 2017). A summary of appeals filed by First Nations is included in Table 5.1.

Balance of Power

In theory, the federal government's decision to approve the TMEP takes precedence over whatever opinions provincial and municipal governments may have on the matter; however, both municipal and provincial governments have permit-granting powers which they can use to their advantage (Hoberg 2013). For example, the City of Burnaby has yet to issue any permits for the TMEP, causing construction delays (Graveland 2017); such delaying tactics are also available to the government of BC (Clogg et al. 2017). This has created tensions among the provincial governments of Alberta and BC, as well as among Alberta and the cities of Burnaby and Vancouver (Graveland 2017).

In a 2013 article in which he develops a framework for political risk analysis applied to pipeline conflicts, Hoberg states that although it is unclear if a province could impede a federally approved project, "provincial authority is considered a "potential" veto point." (Hoberg 2013). A similar argument could be made for municipalities which are in opposition to the TMEP such

Plaintiff	Summary of Argument
Coldwater Indian Band	The Crown neglected its duty to consult based on the lack of proper assessment of the TMEP's impact on the Coldwater Band's main water source.
Squamish Nation	The full impacts of the TMEP on the Squamish Nation's land, water, and resources have not yet been evaluated; the information will not be available until after construction is already underway.
Stk'emlupsemc Te Secwepemc Nation (SSN)	The Crown did not adequately consult the SSN; in particular, the pipeline passes through a sacred site and sensitive grasslands.
Sto:lo	The Crown neglected its duty to consult with respect to the established food, social, and ceremonial fishing rights of the Sto:lo. Issues raised by the Sto:lo during consultation were not adequately addressed.
Tsleil-Waututh Nation	The Crown neglected its duty to consult and unlawfully excluded marine shipping from the environmental assessment.
Upper Nicola Band	The Crown neglected its duty to consult and did not adequately address the Upper Nicola Band's concerns about their rights.

Table 5.1: Summary of arguments presented by First Nations to the Federal Court of Appeal in opposition to the TMEP

as Burnaby and Vancouver, which may yet wield some power in this conflict, and as mentioned in the previous section, First Nations may have the strongest case when it comes to opposing the project. Dealing with and negotiating the balance of power in this conflict could set new precedents for future projects.

5.2.3 Involved Parties

Several parties are involved in this conflict. Some of the key DMs are identified in the next sections.

Trans Mountain

Trans Mountain is committed to constructing the TMEP and still anticipates an in-service date of 2019 (Kinder Morgan Canada Limited 2017b). As of November 2017, the majority of the filing to be done by Trans Mountain regarding the 157 NEB conditions is either under review or has not yet been completed; many of these are due before certain construction activities can begin (National Energy Board 2017b). Limited construction on the TMEP began in September 2017 on terminal facilities (Kinder Morgan Canada Limited 2017a).

In an October 2017 press release, Kinder Morgan President Ian Anderson, stated: "Now, more than ever, this project is critical for our customers and Canada and we remain committed to delivering the project in an environmentally responsible way that respects our extensive and meaningful consultations with Indigenous Peoples, communities and individuals" (Kinder Morgan Canada Limited 2017a). This being said, Trans Mountain is eager to proceed without further delay: according to Kinder Morgan CEO Steve Kean, each month of delay results in $30 to $35 million in costs for the company, and the project is already nine months behind schedule (Kinder Morgan Canada Ltd 2017).

Opposition

First Nations groups, environmentalists, and some municipalities are strong opponents of the TMEP. Lawsuits calling for judicial review of the NEB's initial decision to award the TMEP with approval certificates have been launched by these groups. In October 2017, arguments were presented to the Federal Court of Appeal, with consolidated hearings held for 15 appeals filed by six First Nations groups, the City of Burnaby, BC, the City of Vancouver, BC, and the Raincoast Conservation Foundation and Living Oceans Society. Furthermore, two judicial reviews in the BC Supreme Court regarding the provincial government's issuing of the Environmental Assessment certificate are scheduled for hearings in November 2017.

Federal Government

The federal government is in favour of the TMEP. Apart from approving the NEB's recommendation to move forward with the TMEP, the federal government continues to support the project. Canadian Prime Minister Justin Trudeau reaffirmed the government's support for the TMEP after the provincial elections in BC. Speaking in Rome, Italy, the Prime Minister stated: "The decision we took on the Trans Mountain pipeline was based on facts, evidence on what is in the best interest of Canadians, and indeed all of Canada" (The Canadian Press 2017; Fife and McCarthy 2017). With respect to the impact of BC's election results on the TMEP, Prime Minister Trudeau stated: "Regardless of the change in government in British Columbia or anywhere, the facts and evidence do not change" (The Canadian Press 2017).

Provincial Governments

Although provinces do not have the jurisdictional power to approve or reject pipeline projects, they can grant or reject provincial permits required for the pipeline's construction (Hoberg 2013). The provincial governments of Alberta and BC currently hold opposing views on the TMEP: Alberta supports the project, while BC opposes it.

BC's stance on the TMEP has flip-flopped over time: the Liberal government headed by Premier Christy Clark originally opposed the project and did not give its approval until January 2017, when a fiscal deal was struck between the province and Trans Mountain (Polak and Coleman 2017). Before the May 2017 provincial election, the government of BC continued to support the TMEP.

The provincial election ultimately resulted in a government formed by 41 members of the New Democratic Party (NDP) and 3 members of the Green Party, leaving NDP leader John Horgan as Premier. The two parties signed a Confidence and Supply Agreement, which included a provision to "Immediately employ every tool available to the new government to stop the expansion of the Kinder Morgan pipeline" (B.C. Green Caucus and B.C. New Democrat Caucus 2017). The BC government is currently an intervenor in the consolidated hearings before the Federal Court of Appeal (FCA), arguing that the GIC's approval of the TMEP did not include sufficient justification for the decision.

5.3 Conflict Model

In order to conduct an analysis, the conflict described in the preceding sections needs to be mathematically formalised. This is done using the Option Form which compactly represents a conflict with multiple DMs and options (Howard 1971).

5.3.1 Decision-makers, Options, and Feasible States

In this section, the DMs and options used to construct the conflict model are discussed. Table 5.2 details the DMs, their options, and a description of each option in this conflict.

Trans Mountain

As the corporation responsible for the proposal and construction of the TMEP, Trans Mountain is a DM. The company has several options available. First, it can make economic concessions such as fiscal benefits paid to governments or mutual benefit agreements signed with First Nations in order to make the project more palatable to its opponents. Next, Trans Mountain can engage in time-intensive concessions such as increased consultation with Indigenous groups or undergoing additional environmental assessments. Finally, Trans Mountain can cancel the project altogether.

Government of British Columbia

The BC government can take two broad types of action to oppose the pipeline project. First, BC can try to use its provincial authority to delay or cancel the project. The province could, for example, attach additional conditions to the project beyond those decreed by the NEB, decide to conduct its own environmental assessment of the project, fulfil its constitutional obligation to consult First Nations before granting provincial permits, or, in a more extreme case, reject the approval granted by the Liberal government or reject the project altogether (Clogg et al. 2017).

Second, BC could oppose the project by challenging the federal government's approval of Trans Mountain. This could be done by joining some of the lawsuits currently before the FCA and arguing that the federal government was derelict in its duties to consult with First Nations and to conduct a thorough environmental assessment of the project. Should BC choose neither of the above two options, it will be considered as a tacit approval of the project.

Additional Decision-makers

In addition to Trans Mountain and the BC government, additional decision-makers were considered, but ultimately left out of the conflict model. The reasons for this are twofold: first, as detailed below, the DMs would be limited to pursue a single course of action; second, each additional feasible state results in a large increase of possible rankings (see Section 4.5.2). Based on the storage requirements, a maximum number of 11 feasible states (which would result in 1622632573 possible rankings) could be handled, which in turn required a smaller set of DMs. Based on the progression of the conflict until November 2017, Trans Mountain and the BC government were judged to be the two most important DMs.

Opposition: This group includes First Nations, environmental groups, and municipalities who are opposed to the TMEP. These groups have similar concerns about the pipeline and often work together towards their common goals. These groups are firmly opposed to the project and as such will always deploy legal and other tactics designed to delay or cancel the project.

Support: This group assembles those in favour of the TMEP, including the provincial government of Alberta headed by Andrea Horwath of the NDP and the federal government. Analogous to the opposition group, this set of DMs will always support the project, though their actions are usually limited to voicing public support.

Table 5.2: Decision-makers and options for the Trans Mountain conflict

Decision-maker	Options	Description
Trans Mountain	1. Economic concessions	1: Provide project opponents/critics with economic concessions 0: Do not provide economic concessions
	2. Time-intensive concessions (e.g., increased consultation, additional environmental assessments)	1: Engage in time-intensive activities to appease opponents 0: Do not engage in time-intensive activities to appease opponents
	3. Cancel project	1: Cancel the project 0: Continue with the project
BC Government	1. Opposition at provincial level	1: Oppose the pipeline using provincial authority (e.g., denying permits, adding conditions) 0: Do not oppose the project using provincial authority
	2. Opposition at federal level	1: Oppose the project by challenging the federal government 0: Do not oppose the project by challenging the federal government

5.3.2 Feasible States

Based on the number of options shown in Table 5.2, there are $2^5 = 32$ possible states; however, they are not all feasible. First, note that Trans Mountain cannot simultaneously cancel the project while taking one or more of its remaining actions. Next, the BC government cannot oppose the TMEP at either the provincial or federal level if it has been cancelled. Finally, it is also assumed that it is unlikely that the BC government will oppose the project either at the federal or provincial level if Trans Mountain makes time-intensive concessions, as these would go towards addressing the provincial government's objections. This removes 21 states leaving 11 feasible states in the conflict, which are shown in Option form in Table 5.3 and in graph form in Figure 5.2.

Table 5.3: Feasible states for the Trans Mountain conflict

DM	Options											
Trans Mountain	Economic concessions	0	1	0	1	0	0	1	0	1	0	1
	Time concessions	0	0	1	1	0	0	0	0	0	0	0
	Cancel project	0	0	0	0	1	0	0	0	0	0	0
BC Government	Provincial opposition	0	0	0	0	0	1	1	0	0	1	1
	Federal opposition	0	0	0	0	0	0	0	1	1	1	1
State number		s_0	s_1	s_2	s_3	s_4	s_8	s_9	s_{16}	s_{17}	s_{24}	s_{25}

5.4 Inverse Engineering Methodology

In this section, a brief overview of the ideas underlying the inverse engineering methodology is conducted. The goal of applying the inverse engineering methodology is to provide a DM with useful information for future decision-making. A good starting point is to discuss preference rankings and their role in the process.

5.4.1 Preference Rankings

For the purposes of this analysis, two preference relations exist between two states s_1 and s_2: first, one state can be strictly preferred over the other, which is expressed as $s_1 > s_2$; second, the states can be equally preferred, denoted $s_1 = s_2$. Preference rankings are orderings of states that reflect how the DM rates them with respect to each other. As an example of a preference ranking for the TMEP conflict with the states from Section 5.3.2, consider $s_{25} > s_9 = s_8 > s_{16} > s_4 > s_3 > s_{17} > s_1 > s_0 > s_2 = s_{24}$. This ranking is read left to right from most to least preferred state; it expresses that state s_{25} is the most preferred state and that states s_2 and s_{24} are tied for least preferred state.

Given a preference ranking for states, one can associate an ordinal payoff value to each state. The payoff expresses a state's ranking compared to the others; a state with a higher payoff is more preferred to a state with a lower one. For simplicity, since there are eleven states in the TMEP conflict, the most preferred state receives a payoff of 11, and the payoff decreases by one each time a less preferred state is encountered. The lowest payoff that a state can have is 1. Thus, in the preference ranking given above, state s_{25} has a payoff of 11, states s_9 and s_8 have a payoff of 10, state s_{16} has a payoff of 9, state s_4 has a payoff of 8, and so on until states s_2 and s_{24} which have a payoff of 3. It is important to keep in mind that the payoffs are ordinal and as

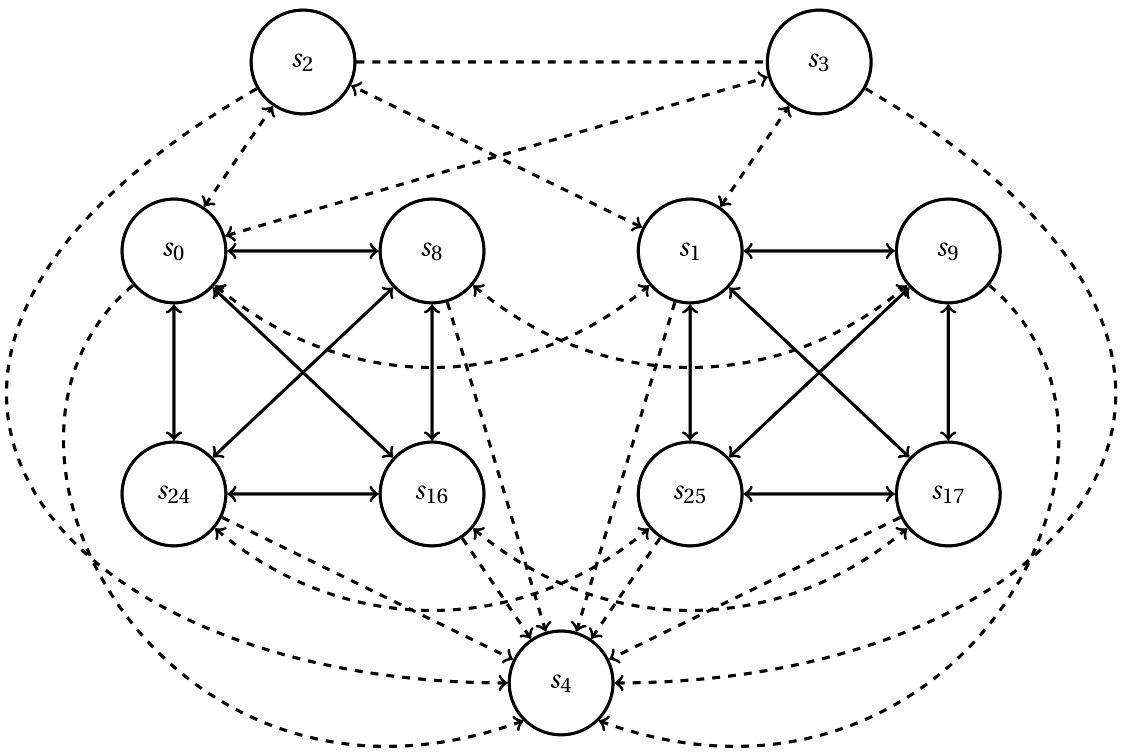

Solid lines: BC goverment
Dashed lines: Trans Mountain

Figure 5.2: Graph Model of the TMEP conflict

such do not express any degree of preference; they simply convey the fact that states with a higher payoff are more preferred to those with lower ones.

Preference rankings are usually determined based on thorough research and/or interviews with DMs; however, these techniques do not guarantee accurate preference rankings. First-hand information regarding the DM's state of mind during the conflict is not always available, and this is particularly true for historical conflicts. Furthermore, DMs may not wish to volunteer information about their preferences to either an analyst or to other DMs involved in the dispute.

The inverse engineering methodology aims to assist analysts by observing a DM's actions and, on that basis, inferring information about the underlying preference ranking which prompted such a move. In the context of the current conflict, ascertaining Trans Mountain's preference could, for example, allow the BC government to determine how likely it is that the project will be cancelled (assuming that the provincial government would like the project to be cancelled). If certain states are less preferred to cancellation, the BC government could steer the conflict in that direction and eventually have the project cancelled.

At the beginning of the conflict, the list of all possible preference rankings is generated for the DM who is the subject of the analysis. As the conflict progresses and observations are logged, this list is narrowed and a clearer picture of that DM's preferences begins to emerge. The process of removing rankings from the list of possible rankings is done with the help of behavioural profiles.

5.4.2 Behavioural Profiles

Each behavioural profile aims to capture how a different type of DM might behave in strategic interactions. Based on the action observed and on the assumed behavioural profile, the preference rankings which are inconsistent with these conditions are excluded from the list of possible rankings. For this case study, the four behavioural profiles described in Chapter 4 (default, Nash, GMR, and SMR) are used to represent four different types of DMs.

5.4.3 Advice Function

The use of preference rankings and behavioural profiles culminates with the ADVICE function. As discussed in Section 4.5, this function provides the observing DM with relevant information based on the evolution of the conflict (Garcia, Obeidi and Hipel 2018).

As will be illustrated in the next section, these results are flexible in that they are designed to accommodate a variety of decision heuristics by the DM conducting the analysis: a DM may be concerned with payoff maximisation, in which case the expected value would guide their decision-making; a DM may worry about sanctioning, in which case the opponent reachable states would be useful information; or a DM may wish to maximise or minimise the probability of a particular state, in which case the next state probabilities are valuable, to name a few examples.

5.5 Analysis Results

Based on the conflict model outlined in Section 5.3, an analysis of the situation was performed using the inverse engineering approach as well as the ADVICE function developed in Chapter 4. The analysis steps and results are discussed in the following sections.

5.5.1 Setup

This analysis is conducted from the point of view of the BC government, who wishes to choose optimal actions based on Trans Mountain's likely preference ranking which is assumed to be unknown to the BC government. Given the new NDP government's statements regarding the TMEP, it is assumed that the government's most preferred state is that in which the TMEP is cancelled. Next, the government would like as many concessions as possible on behalf of Trans Mountain, followed by time concessions, economic concessions, and no concessions. Within its own set of options, the BC government prefers to take federal opposition rather than provincial opposition to the project. The government's preference ranking is as follows: $s_4 \succ s_3 \succ s_2 \succ s_1 \succ s_{17} \succ s_9 \succ s_{25} \succ s_{16} \succ s_8 \succ s_{24} \succ s_0$. The government's payoff for each state is an integer from one to eleven: the highest ranked state receives the highest payoff (state s_4 has a payoff of 11) and each subsequent state's payoff is one less that that of the preceding state (state s_3 has a payoff of 10 and so on until state s_0, which has a payoff of 1).

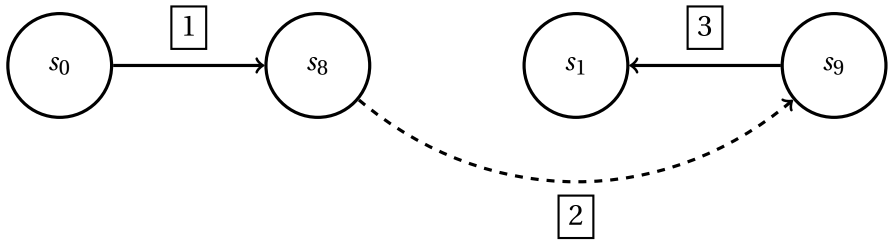

Figure 5.3: Evolution of the Trans Mountain conflict in graph form

5.5.2 Observed Behaviour

In order to apply the inverse engineering approach, it is first necessary to record the DM UMs through the evolution of the conflict. This is done in Table 5.4, which illustrates the progression of the TMEP conflict to the present day and in Figure 5.3. At the start of the conflict, the NEB had just approved the TMEP; none of the DMs had yet made any moves (state s_0). Next, the BC Liberal government voiced its opposition to the project, which was limited to provincial-level tactics (state s_8). Notice that the move from state s_0 to state s_8 constitutes a unilateral move by BC since only BC changes its option choice as indicated by the arrow. In response, Trans Mountain struck a fiscal deal with the Liberal government (state s_9), following which the province changed from opposing the project to supporting it (state s_1). This state remained unchanged until the provincial election (demarcated in the table by a vertical line), after which the NDP/Green government of BC began to oppose the project at the federal level by joining some of the lawsuits against the NEB's original decision (state s_9).

Table 5.4: Evolution of the Trans Mountain conflict

DM	Options					
Trans Mountain	Economic concessions	0	0 $\rightarrow$	1	1	1
	Time concessions	0	0	0	0	0
	Cancel project	0	0	0	0	0
BC Government	Provincial opposition	0 $\rightarrow$	1	1 $\rightarrow$	0 $\rightarrow$	1
	Federal opposition	0	0	0	0	0
State		s_0	s_8	s_9	s_1	s_9

5.5.3 Results and Discussion

Given Trans Mountain's observed moves (from state s_8 to s_9, then staying at state s_1), the ADVICE function is called. Table 5.5 details the number of possible rankings remaining after the algorithms have been run for each behavioural profile. The total number of possible rankings is 1 622 632 573; as shown in Table 5.5, this number is reduced, sometimes significantly, after the ADVICE function is called.

For the Default profile, the logged observations result in 53.15% of the rankings remaining; the highest percentage of any of the behavioural profiles. This is consistent with the fact that the default profile was defined to be the least restrictive: it only uses one criterion for removing rankings. The other behavioural profiles result in a steeper decrease of the number of possible rankings to 10% or less of the original number. Given that the original number of possible rankings is in the order of one billion, there still remain a significant number of rankings (approximately in the order of 100 million) for Nash, GMR, and SMR profiles.

Thus, in contrast to the sustainable development conflict discussed in Chapter 4, the number of possible rankings after the algorithms have been run is too high to produce a tractable list that might be useful to the observing DM. However, as the next sections will illustrate, the information gleaned from the ADVICE function remains highly relevant and valuable.

Table 5.5: Number of rankings remaining after all observations have been entered into the ADVICE function

Behavioural Profile	**Number of possible rankings**	**Percentage of rankings remaining**
Default	862 440 068	53.15%
Nash	139 011 332	8.567%
GMR	98 687 972	6.08%
SMR	163 413 156	10.07%

Default Profile

First, note that from state s_1 (Trans Mountain makes economic concessions, government mounts no opposition), the BC government can expect a payoff of at least 7 if it chooses to stay at state s_1, or move to either state s_9 (provincial opposition) or to state s_{17} (federal opposition). A move to state s_{25} (provincial and federal opposition) would result in a lower expected value at the immediate next step (i.e., when Trans Mountain counter-moves). The expected value information is provided so that the DM has a rough idea of what is likely to occur at the next step; for DMs who prefer to maximise their payoff at each turn, this information is highly valuable.

Looking beyond the expected value, the BC government can also see in more detail how Trans Mountain is likely to respond to its move. For instance, if it chooses to remain at state s_1, it is most likely that Trans Mountain will also remain at that state. In fact, any move by the BC government is most likely to be mirrored by Trans Mountain. Based on this information, the BC government can be somewhat reassured with respect to the possibility of sanctioning by Trans Mountain; since Trans Mountain is likely to stay in whatever state to which the government moves, the threat of sanctioning is diminished.

If the BC government has a particular end result in mind (e.g., cancellation of the TMEP – state s_4), it is also possible to see which of its moves makes this state more likely. Although cancellation is not the most likely response for Trans Mountain at any of the BC government's moves, the highest probability of cancellation occurs if the BC government moves to either state s_9 (25%), state s_{17} (26%), or state s_{25} (26%). On the other hand, if the BC government wishes to avoid a particular state at any cost (e.g., state s_0, its least preferred state), the next state probabilities show that any move away from state s_1 makes state s_0 unreachable for Trans Mountain. The results for the 'default' behavioural profile are summarized in Table 5.6.

State	Expected payoff value	Counter-move probabilities	Trans Mountain's reachable states
s_1	7.83	$s_0 = 17\%$ $s_1 = 32\%$ $s_2 = 17\%$ $s_3 = 17\%$ $s_4 = 17\%$	s_0, s_1, s_2, s_3, s_4
s_9	7.02	$s_4 = 25\%$ $s_8 = 8\%$ $s_9 = 67\%$	s_4, s_8, s_9
s_{17}	7.26	$s_4 = 26\%$ $s_{16} = 26\%$ $s_{17} = 48\%$	s_4, s_{16}, s_{17}
s_{25}	5.77	$s_4 = 26\%$ $s_{24} = 26\%$ $s_{25} = 48\%$	s_4, s_{24}, s_{25}

Table 5.6: Results of the ADVICE function for the BC government after observing Trans Mountain move from state s_8 to state s_9, then staying at state s_1 and assuming a default behavioural profile

Nash Profile

Recall that a DM with a Nash profile is not concerned with risk and has low foresight. DMs who behave according to this profile never move to a less preferred state and remain at states which are Nash stable. Given these assumptions, the results obtained from the ADVICE function are different than those from the previous profile. The results for the Nash behavioural profile are summarized in Table 5.7.

It is now certain that Trans Mountain will also choose to remain at state s_1 and at state s_9 if the BC government does so. However, Trans Mountain's response to the BC government moving to either state s_{17} or to state s_{25} is less certain than it was under the default profile assumption. Although the likelihood of cancellation is decreased compared to the default results, the probabilities that Trans Mountain will choose either of its remaining two options are more equal. Depending on the government's level of risk aversion, they may or may not wish to move to state s_{17} or to state s_{25}.

Again, the information provided can be used in many ways to assist the BC government in its decision-making. Given the BC government's assumed preference ranking, its best move is to remain at state s_1 if the goal is to maximise its payoff. However, if the government would like cancellation to be a possibility, it may opt to move to either state s_{17} or to state s_{25}, taking a short-term disimprovement with the eventual goal of reaching its most preferred state. Note that in this case, the government's least preferred state (state s_0) is not a possibility since it is certain that Trans Mountain will remain at state s_1 and not move to any of the reachable states from there. This is useful for the government to know, since the possibility of sanctioning to state s_0 is no longer a concern.

State	**Expected payoff value**	**Counter-move probabilities**	**Trans Mountain's reachable states**
s_1	8	s_1 = 100%	s_0, s_1, s_2, s_3, s_4
s_9	6	s_9 = 100%	s_4, s_8, s_9
s_{17}	6.53	s_4 = 18% s_{16} = 39% s_{17} = 43%	s_4, s_{16}, s_{17}
s_{25}	4.88	s_4 = 18% s_{24} = 39% s_{25} = 43%	s_4, s_{24}, s_{25}

Table 5.7: Results of the ADVICE function for the BC government after observing Trans Mountain move from state s_8 to state s_9, then staying at state s_1 and assuming a Nash behavioural profile

GMR Profile

DMs who move according to the GMR profile are concerned with possible opponent sanctioning and as such are more risk averse than any of the previous behavioural profiles. Assuming that Trans Mountain is this type of DM yields the results shown in Table 5.8.

Note that moving to state s_{17} (federal opposition) now has the highest expected value for the government, followed by staying at state s_1. The only certainty in this case is that Trans Mountain will remain at state s_1 if the BC government does so. Interestingly, a move away from state s_1 results in a 26% chance of cancellation, regardless of whether the BC government moves to state s_9, s_{17}, or s_{25}; should the government wish to keep the possibility of cancellation alive, it can move to any of its reachable states from state s_1. Also note that states s_8, s_{16}, and s_{24} are unlikely to occur: these are states in which Trans Mountain offers neither economic nor time concessions. In other words, given these results, the BC government does not have to worry about Trans Mountain reversing the economic concessions given to the previous government.

In planning its next move, the BC government could, for example, wish to maximise its expected payoff by moving to state s_{17}. It may prefer certainty and choose to remain at state s_1, or decide that pursuing cancellation is important and move to either state s_9, s_{17}, or s_{25}.

State	Expected payoff value	Counter-move probabilities	Trans Mountain's reachable states
s_1	8	s_1 = 100%	s_0, s_1, s_2, s_3, s_4
s_9	7.31	s_4 = 26% s_9 = 74%	s_4, s_8, s_9
s_{17}	8.05	s_4 = 26% s_{17} = 74%	s_4, s_{16}, s_{17}
s_{25}	6.58	s_4 = 26% s_{25} = 74%	s_4, s_{24}, s_{25}

Table 5.8: Results of the Advice function for the BC government after observing Trans Mountain move from state s_8 to state s_9, then staying at state s_1 and assuming a GMR behavioural profile

SMR Profile

Recall that DMs with an SMR behavioural profile move with three steps in mind: their initial move, their opponent's counter-move, and their response. As such, these DMs have the most foresight out of all of the behavioural profiles and are also risk averse. Assuming that Trans Mountain is this type of DM yields the results shown in Table 5.9.

In this case, the government remaining at state s_1 ensures that Trans Mountain will do the same. Looking at expected values, remaining at state s_1 yields the highest payoff. Should the BC government move to state s_9, Trans Mountain will likely remain at state s_9, but there is a 26% probability that the project will be cancelled. If the government chooses to move to state s_{17} or to state s_{25}, then the likelihood of cancellation of the project increases to 32%. As was the case with the GMR profile, states 8, 16, and 24 have no possibility of occurring; this eases any concerns that the BC government might have about sanctioning.

State	Expected payoff value	Counter-move probabilities	Trans Mountain's reachable states
s_1	8	s_1 = 100%	s_0, s_1, s_2, s_3, s_4
s_9	7.31	s_4 = 26% s_9 = 74%	s_4, s_8, s_9
s_{17}	7.39	s_4 = 32% s_{16} = 30% s_{17} = 38%	s_4, s_{16}, s_{17}
s_{25}	6.03	s_4 = 32% s_{24} = 30% s_{25} = 38%	s_4, s_{24}, s_{25}

Table 5.9: Results of the Advice function for the BC government after observing Trans Mountain move from state s_8 to state s_9, then staying at state s_1 and assuming a Nash behavioural profile

In addition to examining the results for each behavioural profile, it is also useful to compare the results across all of the profiles. In doing so, one can notice that increasing Trans Mountain's risk aversion increases the probability of cancellation. In other words, the more risk-averse Trans Mountain, the more likely it is to cancel the project. The behavioural profile assigned to the opponent should thus be chosen carefully so as to obtain the most relevant results; in this case, the BC government, may, from previous experiences, have an idea of Trans Mountain's level of risk aversion and act accordingly.

5.6 Chapter Summary

By leveraging the information provided by observable opponent moves, the inverse engineering method allows the observing DM (in this case, the BC government) to refine the list of possible opponent's preferences and, most importantly, to use this information to improve their own decision-making process. This is done with the assistance of the ADVICE function, which gives the DM the expected value of each of its possible moves, the probabilities regarding the opponent's counter-move, as well as which states are reachable by the opponent.

With this information in hand, the DM running the analysis (BC government, in this case study) can choose which action is in their best interest. What constitutes the DM's "best interest" is at their discretion; the power of the ADVICE function lies in the fact that it provides the DM with relevant data which the DM can then use to make more informed decisions. For example, a DM may wish to maximise their expected payoff at the immediate next step, to have the highest probability for their most preferred state, to have the lowest probability for their least preferred state, or to take some risk of sanctioning, or follow some other decision heuristic. The output of the ADVICE function is flexible enough to accommodate a variety of decision-making rules and as such is a valuable addition to a DM's toolkit.

Applied to the TMEP conflict, the new methodology provides interesting insights. Under the default profile which has the least restrictive assumptions, the highest expected value occurs at state s_1 (Trans Mountain makes economic concessions), while the highest likelihood of cancellation occurs when the BC government moves to either state s_{17} (Trans Mountain makes economic concessions, BC government opposes at federal level) or to state s_{25} (Trans Mountain makes economic concessions, BC government opposes at federal and provincial levels). If Trans Mountain is assumed to have a Nash profile, state s_1 also has the highest expected value; however, the possibility of cancellation is only conserved by a move to state s_{17} or to state s_{25}. When the GMR profile is assumed, state s_{17} now has the highest expected value and is tied with states s_9 (Trans Mountain makes economic concessions, BC government opposes at provincial level) and s_{25} for highest likelihood of cancellation. Finally, when Trans Mountain's behavioural profile is assumed to be SMR, state s_1 has the highest expected value and the highest probability of cancellation occurs in states s_{17} and s_{25}.

Once again, note that the results obtained are amenable to a variety of decision-making heuristics. The BC government may wish to maximise its expected value, to cancel the project, or to avoid sanctioning, for example; in fact, it is not necessary for the BC government to always follow the same decision-making rules at each call of the function. As its name implies, the ADVICE function's main purpose is to provide relevant information on the basis of which the DM conducting the analysis can make informed decisions.

Future research will explore additional behavioural profiles to the ones used in this case study. The new behavioural profiles could be founded on additional solution concepts from the Graph Model for Conflict Resolution Methodology or based on other decision-making heuristics. Additional output from the ADVICE function will also be explored to further enhance the information it provides to DMs.

Chapter 6

Conclusions and Future Research

The goals of this research are to innovate with respect to mathematical solution concepts and to provide a way forward in the analysis of a conflict in the absence of any preference information from opponents. To these ends, a new family of solution concepts, initial state stabilities, was defined in Chapter 3, while the inverse engineering approach was introduced in Chapters 4 and 5. As will be discussed shortly (Section 6.2.3) these contributions can be nicely integrated.

6.1 Research Contributions

As alluded to above, this work set out to answer two motivating questions:

1. What are other ways to define stability for a state?

2. What can one do if one has no preference information about DMs?

6.1.1 In answer to Q1: What are other ways to define stability for a state?

This question is motivated by a desire to expand the range of DM behaviours that can be mathematically modelled in strategic interactions. As suggested by the literature, capturing different types of behaviours provides DMs and analysts with additional insights into possible resolutions and evolutions of a conflict. The family of initial state stabilities was defined to answer this question: rather than comparing the states which are reachable by opponents from the set of UIs as all of the common Graph Model solution concepts do, comparing reachable states by opponents from the set of UIs as well as from the initial state provides an alternate way to define stability. Chapter 3 delved into how initial state stabilities are defined as well as their applicability to real-world disputes.

The starting point of this chapter is the idea of including opponent reachable states from the starting state in stability determination. This approach is different from that used in the standard Graph Model analysis and allows one to explore the future consequences of remaining at a given starting state. From here, the concept of an initial state, which is then used to define initial state stability, was introduced. Initial state stabilities thus take into account how the conflict evolves both from the focal DM's available UIs *and* from the initial state.

Next, a total of four new solution concepts: pessimistic initial state stability (PIS), optimistic initial state stability (OIS), sequentially pessimistic initial state stability (SPIS), and sequentially optimistic initial state stability (SOIS) were defined for n-DM conflicts. The optimistic solution concepts (OIS, SOIS) model DMs who are interested in comparing best-case scenarios, while the pessimistic solution concepts (PIS, SPIS) compare worst-case outcomes. A further distinction is drawn between PIS/OIS and SPIS/SOIS based on whether the opponent's UMs or UIs are considered.

The latter half of the chapter provided formal proofs relating the family of new solution concepts to existing Graph Model solution concepts and to each other. In special cases, there is equivalence between some of the standard Graph Model concepts and the initial state stabilities as well as equivalence among initial state stabilities. The Elmira conflict was used as a case study to demonstrate the applicability of the new solution concepts as well as the insights they can provide to DMs and analysts.

6.1.2 In answer to Q2: What can one do if one has no preference information about DMs?

This question is motivated by the fact that preferences can be difficult to obtain. This can occur when DMs are not willing to share their preferences or in historical conflicts when there may be a lack of information. As it turns out, meaningful insights and information can be gathered even when DM preferences are unknown. The inverse engineering approach uses a DM's observable actions to infer preference information and to produce useful strategic advice. The theoretical approach introduced in Chapter 4 and demonstrated in Chapter 5 allows for the calculation of expected payoff values and probabilities of counter-moves which can in turn be used by the observing DM to make strategic decisions on how to proceed.

Chapter 4 focused on the introduction of the inverse engineering methodology and its associated algorithms. The key idea behind the methodology is that of leveraging observable behaviour to infer preference information. This idea in turn rests on Revealed Preference Theory and on its heuristic value to model development.

With the idea of using observed behaviour in mind, it is necessary to define behavioural profiles which dictate how the observed DM is making decisions. The behavioural profiles introduced so far are based on Graph Model solution concepts (Nash, GMR, and SMR) and on utility maximisation (Default). These behavioural profiles reflect the types of DMs which are commonly analysed and modelled in conflicts.

Next, the algorithms for the removal of infeasible rankings based on a DM's behavioural profile (Default, Nash, GMR, SMR) and on their observed actions were detailed. Beginning with the list of all possible rankings, the algorithms remove those which are inconsistent with both the observed behaviour and the observed DM's behavioural profile. For smaller conflicts, it is possible to narrow the list of possible rankings to a tractable size; this is not always possible for conflicts with a larger number of states.

The removal of infeasible rankings is taken a step further with the ADVICE function which provides the expected value of the state, the counter-move probabilities, and the opponent reachable states to aid real-time decision-making. This information is useful to the observing DM when they are deciding on their next move as it provides them with a glimpse of what effect their decision might have in the future. As its name implies, the ADVICE function is meant to provide the focal DM with information only, leaving the final decision up to the DM.

At the end of Chapter 4, a simple 2-DM sustainable development conflict was used as a case study. In this conflict, the total number of possible rankings (75) is small enough to result in a shortlist of rankings for the developers.

Chapter 5 applied the ideas and concepts from the previous chapter to the Kinder Morgan pipeline conflict. Conducted from the point of view of the BC government, the conflict started with a total of 1 622 632 573 possible rankings for Trans Mountain. After logging Trans Mountain's observed behaviour and running the ADVICE function for each behavioural profile, this list was greatly reduced. The discussion highlighted the usefulness and flexibility of the data provided by the ADVICE function, which is well suited to a variety of decision-making heuristics.

6.2 Future Research

6.2.1 Bridging the Research Questions

An interesting research avenue is the bridging of the research strands that underlie this work. This can be done by incorporating the four initial state stabilities as behavioural profiles into the inverse engineering methodology. Thus, in addition to the behavioural profiles that have been developed in this work, PIS, SPIS, OIS, and SOIS profiles would also be available. Enriching the set of behavioural profiles will capture a greater variety of human behaviour in conflict situations and widen the applicability of the inverse methodology.

6.2.2 Initial State Stabilities

Initial state stability is a useful addition to the metarational solution concept family. Future research will explore the possibility of operationalising these calculations. Furthermore, although the new solution concepts have here been defined with one opponent response allowed (i.e., looking one move ahead), the underlying idea can, with a few additional definitions and notions, be extended to apply to an arbitrary number of opponent responses. The generalised metarationalities framework (Zeng et al. 2006; Zeng et al. 2007) extends Graph Model solution concepts to an arbitrary number of rounds; a similar generalisation could be applied to these new solution concepts.

6.2.3 Inverse Engineering

Future work in this area can explore and develop additional behavioural profiles. Those presented here are derived from the most commonly used stability concepts from the Graph Model methodology; expanding the range of behavioural profiles is thus an important project.

With regards to the preferences themselves, it may be interesting to explore other preference structures such as those used for uncertain preferences or strength of preferences mentioned in Chapter 4. The current assumption of transitive preferences might also be possible to relax.

Another area of research lies in enhancing the ADVICE function and algorithms in terms of computational complexity and, in the case of the ADVICE function, of its output. Demonstrating and fine-tuning the scalability of the methods and algorithms is an important next step for eventual integration into a DSS. A possible solution to reducing the computational complexity due to the large number of feasible preference rankings is removing certain infeasible rankings even before any moves are observed. This is possible for conflicts in which one has some information regarding the opponents' preferences (e.g. perhaps one state cannot be the most preferred state; any rankings with this characteristic can therefore be removed at the start). Furthermore, in addition to the information that the ADVICE function currently provides to DMs, its output could also include items such as the probability of sanctioning (i.e., how likely it is that the focal DM's move will be sanctioned at the immediate next step).

Extracting useful information from the potentially high number of remaining rankings is also another research avenue. This would be particularly useful for conflicts with a large nunmber of possible rankings such as the one detailed in Chapter 5. It may be, for example, that the remaining rankings have certain characteristics in common (e.g., a particular state is always the most preferred; another state is always ranked 5th or lower, etc.). When this occurs, it would be helpful for DMs and analysts to be able to extract these commonalities from the (long) list of possible rankings. Such information would nicely complement the data from the ADVICE function and assist in decision-making.

The method is also well poised to tie-in with the Graph Model itself as it aids DMs and modellers to develop accurate preference rankings for opponents. A particularly interesting application is the coupling of these methods and algorithms with the inverse approach to the Graph Model which provides DMs and analysts with the preferences required to lead to a particular state (Kinsara, Kilgour and Hipel 2014). DMs and analysts wishing for the conflict to result in a particular state could consult the inverse Graph Model decision support system and obtain the list of preferences that each DM must have to bring this result to fruition; the methods provided in this work would complement this information by verifying whether the generated preferences are feasible given the DMs' behaviour.

Bibliography

Abbeel, P. and A. Y. Ng. 'Apprenticeship learning via inverse reinforcement learning'. In: *Proceedings of the 21st International Conference on Machine Learning (ICML)* (2004), pp. 1–8.

Aljefri, Y., K. W. Hipel and L. Fang. 'General hypergame analysis within the graph model for conflict resolution'. In: *International Journal of Systems Science: Operations & Logistics* (2018), pp. 1–16.

Allan, R. *Submission to Ministerial Panel for Trans Mountain Expansion Project*. http://robynallan.com/wp-content/uploads/2016/10/05/Economist-Robyn-Allans-Submission-To-The-Ministerial-Panel-September-28-2016.pdf. Accessed November 24, 2017. 2016.

An Act to enact the Impact Assessment Act and the Canadian Energy Regulator Act, to amend the Navigation Protection Act and to make consequential amendments to other Acts (First Reading). http://www.parl.ca/DocumentViewer/en/42-1/bill/C-69/first-reading. Accessed February 22, 2018. Feb. 2018.

Ariely, D. and M. I. Norton. 'How actions create – not just reveal – preferences'. In: *Trends in Cognitive Sciences* 12.1 (2008), pp. 13–16.

Bashar, M. A., K. W. Hipel and D. M. Kilgour. 'Fuzzy Preferences in the Graph Model for Conflict Resolution'. In: *IEEE Transactions on Fuzzy Systems* 30.4 (2012), pp. 760–770.

Bashar, M. A., D. M. Kilgour and K. W. Hipel. 'Fuzzy option prioritization for the graph model for conflict resolution'. In: *Fuzzy Sets and Systems* 246 (2014), pp. 34–48.

Bashar, M. A., A. Obeidi, D. M. Kilgour and K. W. Hipel. 'Modeling Fuzzy and Interval Fuzzy Preferences Within a Graph Model Framework'. In: *IEEE Transactions on Fuzzy Systems* 24.4 (4 2016), pp. 765–778. ISSN: 1063-6706. DOI: 10.1109/TFUZZ.2015.2446536.

B.C. Green Caucus and B.C. New Democrat Caucus. *2017 Confidence and Supply Agreement between the BC Green Caucus and the BC New Democrat Caucus*. https://www.bcndp.ca/latest/its-time-new-kind-government-british-columbia. Accessed November 21, 2017. May 2017.

Becklumb, P. *Pipelines: Government Decision-Making*. 2012.

Ben-Haim, Y. and K. W. Hipel. 'The graph model for conflict resolution with information-gap uncertainty in preferences'. In: *Applied Mathematics and Computation* 126 (2–3 2002), pp. 319–340.

Boreal Leadership Council. *Understanding Successful Approaches to Free, Prior, and Informed Consent in Canada. Part I: Recent Developments and Effective Roles for Government, Industry, and Indigenous Communities*. http://borealcouncil.ca/wp-content/uploads/2015/09/BLC_FPIC_Successes_Report_Sept_2015_E.pdf. Accessed November 28, 2017. Sept. 2015.

Boutilier, C. 'A POMDP formulation of preference elicitation problems'. In: *Eighteenth national conference on Artificial intelligence* (2002), pp. 239–246.

Brams, S. J. and D. Wittman. 'Nonmyopic equilibria in 2 x 2 games'. In: *Conflict Management and Peace Science* 6 (1 1981), pp. 39–62.

Bristow, M., L. Fang and K. W. Hipel. 'From Values to Ordinal Preferences for Strategic Governance'. In: *IEEE Transactions on System, Man, and Cybernetics: Systems* 44.10 (2014), pp. 1364–1383.

Bryant, J. W. *Acting Strategically Using Drama Theory*. Boca Raton, Florida: CRC Press, 2015.

Burt, M. *Who Benefits? A Summary of the Economic Impacts That Result From the Trans Mountain Expansion*. http://www.conferenceboard.ca/e-Library/abstract.aspx?did=7648. Accessed November 24, 2017. 2016.

Canadian Environmental Assessment Act S.C. 2012, c.19, s.52. http://laws-lois.justice.gc.ca/PDF/C-15.21.pdf. Accessed February 12, 2018. Dec. 2014.

Clogg, J., E. Kung, G. Smith and A. Gage. *A Legal Toolbox to Defend BC from the Kinder Morgan Trans Mountain Pipeline & Tankers Project*. https://www.wcel.org/publication/legal-toolbox-defend-bc-kinder-morgan-trans-mountain-pipeline-tankers-project. Accessed November 28, 2017. June 2017.

Coastal First Nations Declaration. http://dogwoodinitiative.org/no-tankers/downloadable-files/cfn-declaration-with-release-and-backgrounder. Accessed February 12, 2018. Mar. 2010.

Conlisk, J. 'Why bounded rationality?' In: *Journal of Economic Literature* XXXIV (1996), pp. 669–700.

Earnest, N. K. *Market Prospects and Benefits Analysis of the Trans Mountain Expansion Project for Trans Mountain Pipeline (ULC)*. https://apps.neb-one.gc.ca/REGDOCS/File/Download/2825856. Accessed November 24, 2017. 2015.

Egan, L. C., P. Bloom and L. R. Santos. 'Choice-induced Preferences in the Absence of Choice: Evidence from a Blind Two Choice Paradigm with Young Children and Capuchin Monkeys'. In: *Journal of Experimental Social Psychology* 46 (1 1959), pp. 204–207.

Environmental Assessment Act SBC 2002, c. 43. http://www.bclaws.ca/Recon/document/ID/freeside/00_02043_01. Current to November 1, 2017. Accessed November 20, 2017. May 2002.

Fang, L., K. W. Hipel, D. M. Kilgour and X. Peng. 'A decision support system for interactive decision making - Part II: analysis and output interpretation'. In: *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 33.1 (2003), pp. 56–66. DOI: 10.1109/TSMCC.2003.809360.

Fang, L., K. W. Hipel, D. M. Kilgour and X. Peng. 'A decision support system for interactive decision making-Part I: model formulation'. In: *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 33.1 (2003), pp. 42–55. DOI: 10.1109/TSMCC.2003.809361.

Fang, L., K. W. Hipel and D. M. Kilgour. *Interactive decision making: The graph model for conflict resolution*. New York: J. Wiley, 1993.

Festinger, L. and J. M. Carlsmith. 'Cognitive consequences of forced compliance'. In: *Journal of Abnormal and Social Psychology* 58 (1959), pp. 203–211.

Fife, R. and S. McCarthy. *Trudeau defends Trans Mountain as BC Greens, NDP gear up for fight.* https://www.theglobeandmail.com/news/politics/trudeau-trans-mountain-pipeline-bc/article35151256/. Accessed November 21, 2017. May 2017.

Figueira, J. R., S. Greco, B. Roy and R. Słowiński. 'An Overview of ELECTRE Methods and their Recent Extensions'. In: *Journal of Multi-Criteria Decision Analysis* 20.1-2 (2013), pp. 61–85. ISSN: 1099-1360. DOI: 10.1002/mcda.1482.

Fraser, N. and K. W. Hipel. *Conflict analysis: models and resolution.* New York: North Holland, 1984.

Fraser, N. and K. W. Hipel. 'Solving complex conflicts'. In: *IEEE Transactions on Systems, Man and Cybernetics* 9 (12 1979), pp. 805–816.

Galeotti, A., S. Goyal, M. O. Jackson, F. Vega-Redondo and L. Yariv. 'Network Games'. In: *The Review of Economic Studies* 77.1 (2010), p. 218. DOI: 10.1111/j.1467-937X.2009.00570.x.

Garcia, A. and K. W. Hipel. 'Inverse engineering preferences in simple games'. In: *Applied Mathematics and Computation* 311 (2017), pp. 184–194. ISSN: 0096-3003. DOI: https://doi.org/10.1016/j.amc.2017.05.016. URL: http://www.sciencedirect.com/science/article/pii/S0096300317303041.

Garcia, A., A. Obeidi and K. W. Hipel. 'Initial State Stability for n-Decision-Maker Conflicts'. In: *INFOR: Information Systems and Operational Research* (). Accepted, to be published in a forthcoming issue.

Garcia, A., A. Obeidi and K. W. Hipel. 'Strategic Advice for Decision-making Under Conflict Based on Observed Behaviour'. In: *Applied Mathematics and Computation* 332 (2018), pp. 96–104. DOI: ttps://doi.org/10.1016/j.amc.2018.03.031.

Garcia, A., A. Obeidi and K. W. Hipel. 'Two methodological perspectives on the Energy East Pipeline conflict'. In: *Energy Policy* 91 (2016), pp. 397–409.

Good, I. J. 'The number of orderings of n candidates when ties are permitted'. In: *Fibonacci Quarterly* 13 (1975), pp. 11–18.

Government of British Columbia. *Trans Mountain Expansion Project Schedule B: Table of Conditions for an Environmental Assessment Certificate.* https://projects.eao.gov.bc.ca/api/document/5892318fb637cc02bea16484/fetch. Accessed November 20, 2017. Jan. 2017.

Government of Canada. *National Energy Board Act, R.S.C., 1985, c. N-7.* Oct. 2014.

Governor in Council. *PC 2016-1069, November 29, 2016.* http://www.gazette.gc.ca/rp-pr/p1/2016/2016-12-10/html/sup1-eng.php. Accessed November 21, 2017. Dec. 2016.

Graveland, B. *'Overreaching extensively:' Notley wants Trans Mountain delays dealt with.* http://www.cbc.ca/news/canada/calgary/alberta-kinder-morgan-notley-pipeline-1.4387475. Accessed November 28, 2017. Nov. 2017.

Gross, O. A. 'Preferential Arrangements'. In: *The American Mathematical Monthly* 69 (1 1975), pp. 4–8.

Hamouda, L., D. M. Kilgour and K. W. Hipel. 'Strength of preference in graph models for multiple decision maker conflicts'. In: *Applied Mathematics and Computation* 179 (2006), pp. 314–327.

Hamouda, L., D. M. Kilgour and K. W. Hipel. 'Strength of Preference in the Graph Model for Conflict Resolution'. In: *Group Decision and Negotiation* 13.5 (2004), pp. 449–462. DOI: 10.1023/B:GRUP.0000045751.21207.35.

Han, X., T. D. L. Nguyen and H. Xu. ‘Conflict analysis based on the graph model for conflict resolution and grey matrix’. In: *Grey Systems: Theory and Application* 3.1 (2013), pp. 95–106. DOI: 10.1108/20439371311293723.

Hands, D. W. ‘Foundations of Contemporary Revealed Preference Theory’. In: *Erkenntnis* 78 (5 2013), pp. 1081–1108.

Hashimoto, K. ‘Multigame effect in finite populations induces strategy linkage between two games’. In: *Journal of Theoretical Biology* 345 (2014), pp. 70–77. DOI: http://doi.org/10.1016/j.jtbi.2013.12.007.

Hausman, D. M. ‘Revealed preference, belief, and game theory’. In: *Economics and Philosophy* 16 (2000), pp. 99–115.

He, S., D. M. Kilgour, K. W. Hipel and M. A. Bashar. ‘A Basic Hierarchical Graph Model for Conflict Resolution with Application to Water Diversion Conflicts in China’. In: *INFOR* 51.3 (Aug. 2013), pp. 103–119. URL: http://search.proquest.com.proxy.lib.uwaterloo.ca/docview/1562515173?accountid=14906.

He, S., D. M. Kilgour and K. W. Hipel. ‘A general hierarchical graph model for conflict resolution with application to greenhouse gas emission disputes between USA and China’. In: *European Journal of Operational Research* 257 (3 2016), pp. 919–932. DOI: 10.1016/j.ejor.2016.08.014.

Hipel, K. W. ‘Conflict Resolution’. In: *Our Fragile World: Challenges and Opportunities for Sustainable Development (Forerunner to the Encyclopedia of Life Support Systems)*. Ed. by M. Tolba. Vol. 1. Oxford, United Kingdom: Eolss Publishers, 2001, pp. 935–952.

Hipel, K. W., L. Fang, D. M. Kilgour and M. Haight. ‘Environmental conflict resolution using the graph model’. In: *Proceedings of the 1993 IEEE international conference on systems, man, and cybernetics* (1993), pp. 153–158.

Hipel, K. W., D. M. Kilgour and R. Kinsara. ‘Strategic Investigations of Water Conflicts in the Middle East’. In: *Group Decision and Negotiation* 45.3 (5 May 2014), pp. 355–376.

Hoberg, G. ‘The Battle Over Oil Sands Access to Tidewater: A Political Risk Analysis of Pipeline Alternatives’. In: *Canadian Public Policy* 39 (3 Sept. 2013), pp. 371–392. DOI: 10.3138/CPP.39.3.371.

Hofbauer, J. and K. Sigmund. ‘Evolutionary Game Dynamics’. In: *Bulletin of the American Mathematical Society* 40 (4 2003), pp. 479–519.

Howard, N. *Confrontation Analysis: How to win operations other than war*. Washington, D.C.: CCRP Publications, 1999.

Howard, N. *Paradoxes of Rationality: Theory of Metagames and Political Behaviour*. MIT Press, 1971.

Hughes, J. D. *Will the Trans Mountain Pipeline and Tidewater Access Boost Prices and Save Canada’s Oil Industry?* https://www.policyalternatives.ca/tidewater-access. Accessed November 24, 2017. 2017.

Indigenous Nations of the Fraser River Watershed. *Save The Fraser Declaration*. Dec. 2010.

Inohara, T. ‘State transition time analysis in the Graph Model for Conflict Resolution’. In: *Applied Mathematics and Computation* 274 (2016), pp. 372–382.

Inohara, T. and K. W. Hipel. ‘"Coalition Analysis in the Graph Model for Conflict Resolution"’. In: *Systems Engineering* 11.4 (2008), pp. 343–359. ISSN: 1520-6858. DOI: 10.1002/sys.20104.

International Treaty to Protect the Salish Sea. https://twnsacredtrust.ca/coast-salish-nations-sign-international-treaty-to-protect-the-salish-sea/. Accessed November 28, 2017. Sept. 2014.

Ke, G., B. Fu, M. De and K. W. Hipel. 'A Hierarchical Multiple Criteria Model for Eliciting Relative Preferences in Conflict Situations'. In: *Journal of Systems Science and Systems Engineering* 21.1 (2012), pp. 56–76.

Ke, G., K. W. Li and K. W. Hipel. 'An integrated multiple criteria preference ranking approach to the Canadian west coast port congestion conflict'. In: *Expert Systems with Applications* 39.10 (2012), pp. 9181–9190. DOI: https://doi.org/10.1016/j.eswa.2012.02.086.

Ke, G., K. W. Li and K. W. Hipel. 'The integration of a modified AHP preference ranking approach and the graph model analysis'. In: *Proceedings of the 2007 IEEE international conference on Systems, man, and cybernetics* (2007), pp. 1830–1835.

Keeney, R. *Value-focused thinking: a path to creative decision making*. Cambridge, Massachussets: Harvard University Press, 1992.

Kilgour, D. M. and C. Eden. *Handbook of group decision and negotiation*. New York: Springer Dordrecht, 2010.

Kilgour, D. M., K. W. Hipel and L. Fang. 'The graph model for conflicts'. In: *Automatica* 23.1 (1987), pp. 41–55.

Kilgour, D. M., K. W. Hipel, L. Fang and X. Peng. '"Coalition Analysis in Group Decision Support"'. In: *Group Decison and Negotiation* 10.2 (2001), pp. 159–175. DOI: 10.1023/A:1008713120075.

Kinder Morgan Canada Limited. *Kinder Morgan Canada Limited Declares Dividends and Announces Results for Third Quarter of 2017*. https://www.transmountain.com/news/2017/kinder-morgan-canada-limited-declares-dividends-and-announces-results-for-third-quarter-of-2017. Accessed November 21, 2017. Oct. 2017.

Kinder Morgan Canada Limited. *Project Benefits Overview: Trans Moutain Expansion Project*. https://www.transmountain.com/project-overview. Accessed November 17, 2017. Nov. 2017.

Kinder Morgan Canada Ltd. *Kinder Morgan Canada Limited Declares Dividends and Announces Results for Thrid Quarter of 2017*. http://ir.kindermorgan.com/annual-quarterly-reports. Accessed November 28, 2017. Oct. 2017.

Kinsara, R., D. M. Kilgour and K. W. Hipel. 'Inverse Approach to the Graph Model for Conflict Resolution'. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45.5 (5 May 2014), pp. 734–742. ISSN: 2168-2216. DOI: 10.1109/TSMC.2014.2376473.

Kinsara, R., O. Petersons, K. W. Hipel and D. M. Kilgour. 'Advanced Decision Support System for the Graph Model for Conflict Resolution'. In: *Journal of Decision Systems, Special Issue on Integrated Decision Support Systems* 24.2 (2015), pp. 117–145. DOI: 10.1080/12460125.2015.1046682.

Kuang, H., M. A. Bashar, K. W. Hipel and D. M. Kilgour. 'Grey-Based Preference in a Graph Model for Conflict Resolution With Multiple Decision Makers'. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45 (9 2015), pp. 1254–1267.

Kung, E. *Saw you in court: The Kinder Morgan Federal Court of Appeal hearing (explained in road signs)*. https://www.wcel.org/blog/saw-you-in-court-kinder-morgan-federal-court-appeal-hearing-explained-in-road-signs. Accessed November 21, 2017. Oct. 2017.

Li, K. W., K. W. Hipel, D. M. Kilgour and L. Fang. 'Preference Uncertainty in the Graph Model for Conflict Resolution'. In: *IEEE Transactions on Systems, Man, and Cybernetics Part A* 34.4 (2004), pp. 507–520.

Li, K. W., T. Inohara and H. Xu. 'Coalition analysis with preference uncertainty in group decision support'. In: *Applied Mathematics and Computation* 231 (2014), pp. 307–319.

Li, K. W., D. M. Kilgour and K. W. Hipel. 'Status Quo Analysis in the Graph Model for Conflict Resolution'. In: *Journal of the Operational Research Society* 56 (2005), pp. 699–707.

Madani, K. 'Modeling international climate change negotiations more responsibly: Can highly simplified game theory models provide reliable policy insights?' In: *Ecological Economics* 90 (2013), pp. 68–76. DOI: 10.1016/j.ecolecon.2013.02.011.

Madani, K. and K. W. Hipel. 'Non-cooperative stability definitions for strategic analysis of generic water resources conflict'. In: *Journal of Water Resources Management* 25 (2011), pp. 1949–1977.

Matbouli, Y., K. W. Hipel and D. M. Kilgour. 'Strategic analysis of the Great Canadian Hydroelectric Power Conflict'. In: *Energy Strategy Reviews* 4 (2014), pp. 43–51.

Matbouli, Y., D. M. Kilgour and K. W. Hipel. 'Robustness of equilibria in the graph model for conflict resolution'. In: *Journal of Systems Science and Systems Engineering* 24.4 (2015), pp. 450–465.

National Energy Board. *Canadian Energy Dynamics Review of 2014*. Feb. 2015.

National Energy Board. *Canadian Energy Overview 2013 - Energy Briefing Note*. https://www.neb-one.gc.ca/nrg/ntgrtd/mrkt/vrvw/2013/index-eng.html. June 2014.

National Energy Board. *Conditions and Requirements: Trans Mountain Pipeline ULC - Trans Mountain Pipeline Expansion*. http://www.neb-one.gc.ca/pplctnflng/mjrpp/trnsmntnxpnsn/index-eng.html#details-panel2. Accessed November 15, 2017. 2017.

National Energy Board. *Crude Oil and Petroleum Products*. https://www.neb-one.gc.ca/nrg/sttstc/crdlndptrlmprdct/index-eng.html. May 2015.

National Energy Board. *National Energy Board Report Trans Mountain Expansion Project (OH-001-2014)*. 2016.

National Energy Board. *NEB's Full Lifecycle Oversight*. http://www.neb-one.gc.ca/bts/nws/rgltrsnpshts/2016/09rgltrsnpsht-eng.html. Accessed November 15, 2017. 2016.

National Energy Board. *NEB's Regulatory Framework*. http://www.neb-one.gc.ca/bts/nws/rgltrsnpshts/2016/08rgltrsnpsht-eng.html. Accessed November 15, 2017. 2016.

National Energy Board. *Trans Mountain interactive conditions tool*. https://www.neb-one.gc.ca/pplctnflng/mjrpp/trnsmntnxpnsn/cndtnstl-eng.html. Accessed November 21, 2017. Nov. 2017.

National Energy Board. *Trans Mountain Pipeline ULC - Trans Mountain Expansion*. http://www.neb-one.gc.ca/pplctnflng/mjrpp/trnsmntnxpnsn/index-eng.html#details-panel3. Accessed November 15, 2017. 2017.

Natural Resources Canada. *About Crude Oil and Petroleum Products*. http://www.nrcan.gc.ca/energy/fuel-prices/4597. Sept. 2014.

Natural Resources Canada. *Oil Sands: A strategic resource for Canada, North America and the global market*. Feb. 2013. URL: http://www.nrcan.gc.ca/energy/publications/6539.

Natural Resources Canada. *Trans Mountain Expansion Project: Project Summary*. http://www.nrcan.gc.ca/energy/resources/19142. Accessed November 15, 2017. 2017.

Ng, A. Y. and S. Russell. ‘Algorithms for inverse reinforcement learning’. In: *Proceedings of the 17th International Conference on Machine Learning (ICML)* (2000), pp. 663–670.

Nosek, G. ‘Re-Imagining Indigenous Peoples’ Role in Natural Resource Development Decision Making; Implementing Free, Prior and Informed Consent in Canada through Indigenous Legal Traditions’. In: *U.B.C. Law Review* 50 (1 Feb. 2017). Accessed November 28, 2017, pp. 95–160.

Nowak, M. A., S. Bonhoeffer and R. M. May. ‘More Spatial Games’. In: *International Journal of Bifurcation and Chaos* 04.01 (1994), pp. 33–56. DOI: 10.1142/S0218127494000046.

Obeidi, A., K. W. Hipel and D. M. Kilgour. ‘The role of emotions in envisioning outcomes in conflict analysis’. In: *Group Decision and Negotiation* 14 (2005), pp. 184–200.

Obeidi, A., D. M. Kilgour and K. W. Hipel. ‘Perceptual Graph Model Systems’. In: *Group Decision and Negotiation* 18 (2005), pp. 261–277.

Office of the Prime Minister. *Prime Minister Justin Trudeau’s Pipeline Announcement*. https://pm.gc.ca/eng/news/2016/11/30/prime-minister-justin-trudeaus-pipeline-announcement. Accessed November 15, 2017. 2016.

Perc, M. ‘The Matthew effect in empirical data’. In: *J. R. Soc. Interface* 11 (98 2014).

Perc, M., J. Gomez-Gardenes, A. Szolnoki, L. M. Floria and Y. Moreno. ‘Evolutionary dynamics of group interactions on structured populations: a review’. In: *Journal of the Royal Society Interface* 10 (20120997. 2013). DOI: http://dx.doi.org/10.1098/rsif.2012.0997.

Philpot, S., K. W. Hipel and P. A. Johnson. ‘Design and Application of a Values-Focused Decision Support System to Facilitate Public Participation in Environmental Decision-Making’. In: *Proceedings of the 17th International Conference on Group Decision and Negotiation* (2017). Ed. by M. Schoop and D. Kilgour.

Philpot, S., K. W. Hipel and P. A. Johnson. ‘Strategic analysis of a water rights conflict in the south western United States’. In: *Journal of Environmental Management* 180 (2016), pp. 247–256.

Polak, M. and R. Coleman. *Reasons for Ministers’ Decision: Trans Mountain Expansion Project*. http://a100.gov.bc.ca/appsdata/epic/html/deploy/epic_project_doc_list_459_a_waa.html. Accessed November 20, 2017. Jan. 2017.

Rego, L. C. and A. M. dos Santos. ‘Probabilistic Preferences in the Graph Model for Conflict Resolution’. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45 (4 2015), pp. 595–608.

Rego, L. C. and G. Vieira. ‘Symmetric Sequential Stability in the Graph Model for Conflict Resolution with Multiple Decision Makers’. In: *Group Decision and Negotiation* 26 (4 2017), pp. 775–792.

Roca, C., J. Cuesta and A. Sanches. ‘Evolutionary game theory: temporal and spatial effects beyond replicator dynamics’. In: *Physics of Life Reviews* 6 (2009), pp. 208–249. DOI: doi:10.1016/j.plrev.2009.08.001.

Roy, B. ‘Classement et choix en présence de points de vue multiples (la méthode ELECTRE)’. In: *La Revue d’Informatique et de Recherche Opérationelle (RIRO)* 8 (1968), pp. 57–75.

Rubinstein, A. and Y. Salant. ‘Some Thoughts on the Principle of Revealed Preference’. In: *The Foundations of Positive and Normative Economics: A Handbook*. Ed. by A. Caplin and A. Schotter. New York: Oxford Univeristy Press, 2008, pp. 115–124.

Saaty, R. W. ‘The analytic hierarchy process—what it is and how it is used’. In: *Mathematical Modelling* 9.3 (1987), pp. 161–176. DOI: https://doi.org/10.1016/0270-0255(87)90473-8.

Saaty, T. L. *The Analytic Hierarchy Process*. New York: McGraw-Hill, 1980.

Sakakibara, H., N. Okada and D. Nakase. 'The Application of Robustness Analysis to the Conflict With Incomplete Information'. In: *IEEE Transactions of Systems, Man and Cybernetics Part C: Applications and Reviews* 32 (1 2002), pp. 14–23.

Samuelson, P. A. 'A Note on the Pure Theory of Consumer's Behaviour'. In: *Economica* 5.17 (1938), pp. 61–71.

Samuelson, P. A. 'Consumption Theory in Terms of Revealed Preference'. In: *Economica* 15.60 (1948), pp. 243–253.

Sen, A. 'Behaviour and the Concept of Preference'. In: *Economica* 40.159 (1973), pp. 241–259.

Sen, A. 'Internal consistency of choice'. In: *Econometrica* 61 (1993), pp. 495–521.

Sharot, T., B. D. Martino and R. J. Dolan. 'How Choice Reveals and Shapes Expected Hedonic Outcome'. In: *Journal of Neuroscience* 29 (12 2009), pp. 3760–3765.

Simon, H. A. *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. New York: Wiley, 1957.

Species at Risk Act S.C. 2002, c. 29. http://laws-lois.justice.gc.ca/PDF/S-15.3.pdf. Accessed February 12, 2018. Nov. 2014.

Spiegler, R. 'On two Points of View Regarding revealed Preference and behavioral Economics'. In: *The Foundations of Positive and Normative Economics: A Handbook*. Ed. by A. Caplin and A. Schotter. New York: Oxford Univeristy Press, 2008, pp. 95–115.

Szabo, G. and G. Fath. 'Evolutionary games on graphs'. In: *Physics Reports* (446 2007), pp. 97–216. DOI: doi:10.1016/j.physrep.2007.04.004.

Tasker, J. P. *Ottawa to scrap National Energy Board, overhaul environmental assessment process for major projects*. http://www.cbc.ca/news/politics/liberal-environmental-assessment-changes-1.4525666. Accessed February 22, 2018. Feb. 2018.

The Canadian Press. *Kinder Morgan Canada shares fall in debut amid political uncertainty in B.C.* http://www.cbc.ca/news/business/kinder-morgan-ipo-morning-1.4137439. Accessed November 21, 2017. May 2017.

Trans Mountain. *Project Overview*. https://www.transmountain.com/project-overview. Accessed November 15, 2017. 2017.

Trans Mountain Pipeline ULC. *Trans Mountain Expansion Project: An Application Pursuant to Section 52 of the National Energy Board Act, Volume 1: Summary*. Dec. 2013.

Treaty Alliance Against Tar Sands Expansion. http://www.treatyalliance.org/treaty/. Accessed November 28, 2017. Sept. 2015.

von Neumann, J. and O. Morgenstern. *Theory of Games and Economic Behaviour*. Princeton, New Jersey, USA: Princeton University Press, 1944.

Walker, S. B., K. W. Hipel and T. Inohara. 'Strategic decision making for improved environmental security: coalitions and attitude'. In: *Journal of Systems Science and Systems Engineering* 18.4 (2009), pp. 461–476.

Walker, S. B., K. W. Hipel and H. Xu. 'A Matrix Representation of Attitudes in Conflicts'. In: *IEEE Transactions on System, Man, and Cybernetics: Systems* 43.6 (2013), pp. 1328–1342.

Walker, S., K. W. Hipel and T. Inohara. 'Dominating attitudes in the graph model for conflict resolution'. In: *Journal of Systems Science and Systems Engineering* 241 (3 2012), pp. 316–336.

Wang, J., K. W. Hipel, L. Fang and Y. Dang. 'Matrix Representations of the Inverse Problem in the Graph Model for Conflict Resolution'. In: *European Journal of Operational Research* 270 (1 2018), pp. 282–293.

Wardil, L. and J. K. da Silva. 'The evolution of cooperation in mixed games'. In: *Chaos, Solitons & Fractals* 56 (2013), pp. 160–165. DOI: 10.1016/j.chaos.2013.07.018.

Xu, H., K. W. Hipel and D. M. Kilgour. 'Matrix Representation of Solution Concepts in Multiple Decision Maker Graph Models'. In: *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 39 (1 2009), pp. 96–108.

Xu, H., K. W. Hipel and D. M. Kilgour. 'Multiple levels of preference in interactive strategic decisions'. In: *Discrete Applied Mathematics* 57 (2009), pp. 3300–3313.

Xu, H., K. W. Hipel, D. M. Kilgour and Y. Chen. 'Combining strength and uncertainty for preferences in the graph model for conflict resolution with multiple decision makers'. In: *Theory and Decision* 69.4 (2010), pp. 497–521.

Xu, H., K. W. Hipel, D. M. Kilgour and L. Fang. *Conflict Resolution Using the Graph Model: Strategic Interactions in Competition and Cooperation*. Studies in Systems, Decision and Control, Springer, Chams, Switzerland. 2018.

Xu, H., D. M. Kilgour, K. W. Hipel and E. A. McBean. 'Theory and implementation of coalitional analysis in cooperative decision making'. In: *Theory and Decision* 76 (2 2014), pp. 147–171.

Xu, H., K. W. Li, K. W. and D. M. Kilgour. 'A matrix approach to status quo analysis in the graph model for conflict resolution'. In: *Applied Mathematics and Computation* 212.2 (2009), pp. 470–480. DOI: 10.1016/j.amc.2009.02.051. URL: http://dx.doi.org/10.1016/j.amc.2009.02.051.

Yu, J., D. M. Kilgour, K. W. Hipel and M. Zhao. 'Power asymmetry in conflict resolution with application to a water pollution dispute in China'. In: *Water Resources Research* 51.10 (2015), pp. 8627–8645.

Zeng, D.-Z., L. Fang, K. W. Hipel and D. M. Kilgour. 'Generalized metarationalities in the graph model for conflict resolution'. In: *Discrete Applied Mathematics* 154.16 (2006), pp. 2430–2443.

Zeng, D.-Z., L. Fang, K. W. Hipel and D. M. Kilgour. 'Policy Equilibrium and Generalized Metarationalities for Multiple Decision-Maker Conflicts'. In: *IEEE Transactions on Systems, Man and Cybernetics* 37.4 (2007), pp. 456–463.

Zeng, D.-Z., L. Fang, K. W. Hipel and D. M. Kilgour. 'Policy Stable States in the Graph Model for Conflict Resolution'. In: *Theory and Decision* 57.4 (2005), pp. 345–365.